DOCUMENT RESUME

ED 205 944 CS 006 264

AUTHOR	Lansman, Marcy, Ed.; Hunt, Earl, Ed.
TITLE	Proceedings of the Lake Wilderness Attention Conference. Interim Technical Report, August 1, 1980 through September 30, 1980.
INSTITUTION	Washington Univ., Seattle. Dept. of Psychology.
SPONS AGENCY	Office of Naval Research, Washington, D.C. Psychological Sciences Div.
REPORT NO	ONR-TR-10
PUB DATE	10 Jul 81
CONTRACT	N00014-80-C-0631
NOTE	373p.
EDRS PRICE	MF01/PC15 Plus Postage.
DESCRIPTORS	*Attention; Attention Control; Behavioral Science Research; *Cognitive Processes; Models; Motivation; Orientation; Psychological Studies; *Reading Research; *Theories; *Visual Perception

ABSTRACT

This technical report contains papers prepared by the 11 speakers at the 1980 Lake Wilderness (Seattle, Washington) Conference on Attention. The papers are divided into general models, physiological evidence, and visual attention categories. Topics of the papers include the following: (1) willed versus automatic control of behavior; (2) multiple resources in time-sharing; (3) a computational theory of thought applied to traditional problems; (4) brain mechanisms in human information processing, and implications for a cognitive theory of processing resources; (5) electrophysiological studies of attention; (6) attention and automaticity; (7) the role of attention in object perception; (8) consequences of visual orienting; (9) unitization and automaticity in reading; (10) a comparison of two-state and continuous-state attention models; and (11) automatism and attentional processes. (RL)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED205944

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

Proceedings
# of the
Lake Wilderness Attention Conference

Marcy Lansman and Earl Hunt, Editors
University of Washington

Technical Report No. 10
July 10, 1981

This research was sponsored by:

Personnel and Training Research Programs
Psychological Sciences Divisions
Office of Naval Research
Under Contract No. N00014-80-C-0631
Contract Authority Identification Number 150-457

Approved for public release distribution unlimited

Reproduction in whole or in part is permitted
for any purpose of the U. S. Government.

CS806264

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report #10	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Proceedings of the Lake Wilderness Attention Conference		5. TYPE OF REPORT & PERIOD COVERED Interim August 1, 1980-Sept. 30, 1980
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Marcy Lansman & Earl Hunt, Editors		8. CONTRACT OR GRANT NUMBER(s) N00014-80-C-0631
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Psychology NI-25 University of Washington Seattle, Washington 98195		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61153N RR042-04; RR042-04-01 NR 150-457
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Psychological Sciences Divisions Office of Naval Research (Code 458) Arlington, Virginia 22217		12. REPORT DATE July 10, 1981
		13. NUMBER OF PAGES 364
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release, distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Attention, automaticity, visual attention, visual search, models of attention, visual perception, evoked potential, reaction time.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report contains papers by the eleven speakers at the Lake Wilderness Attention Conference held in Seattle, Washington, Septemper 22-24, 1980, and jointly sponsored by the Office of Navel Research and the Air Force Office of Scientific Research. The papers have been grouped into three categories; General Models, Physiological Evidence, and Visual Attention.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

CONTENTS

Preface i

Part One: General Models

1. Attention to Action: Willed and Automatic Control of Behavior 1
By Donald A. Norman

2. Multiple Resources in Time-Sharing 32
By Chris Wickens

3. A Computational Theory of Thought Applied to Traditional Problems 53
By Earl Hunt

Part Two: Physiological Evidence

4. Brain Mechanisms in Human Information Processing: Implications for a Cognitive Theory of Processing Resources 108
By Jackson Beatty

5. Electrophysiological Studies of Attention 167
By Steven A. Hillyard

Part Three: Visual Attention

6. Remarks on Attention and Automaticity 194
By Daniel Kahneman and Anne Treisman

7. The Role of Attention in Object Perception 225
By Anne Treisman

8. Consequences of Visual Orienting 250
By Michael I. Posner and Yoav Cohen

9. Unitization and Automaticity in Reading 262
By David LaBerge

10. A Test Between Two-State and Continuous State Attention Models 272
By Marilyn L. Shaw and Robert Mulligan

11. Automatism and Attentional Processes 333
By Richard Shiffrin

Distribution List

Preface

In September, 1980, the Office of Naval Research and Air Force Office of Scientific Research jointly sponsored the Lake Wilderness Conference on Attention. The conference was hosted by the University of Washington at their Lake Wilderness facility outside of Seattle. This technical report includes papers by the eleven speakers at that conference. In some cases the papers are revised versions of the actual transcripts of the talks. In other cases, the authors have chosen to entirely rewrite their presentations. We have reproduced the papers with a minimum of editing and retyping in order to shorten the period between submission and circulation.

Attention to Action: Willed and Automatic Control of Behavior

Donald A. Norman

Program in Cognitive Science
Center for Human Information Processing
University of California, San Diego
La Jolla, California 92093

Tim Shallice

Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge, CB2 2EF England

- Abstract
- The Structure of the Paper
  - Automatic performance
  - Relationship to previous work
- A Theoretical Framework for the Analysis of Attention
  - Schemas and processing structures
  - Horizontal threads
  - Schema selection mechanisms
  - Vertical threads
    - Attentional resources.
    - Motivation.
  - Contention scheduling
    - Trigger conditions.
    - The selection mechanism.
    - Fine timing of operations.
- The Interpretation of the Phenomena
  - The Experimental literature
    - Simultaneous tasks.
    - Attentional demands at the initiation of actions.
  - Experiential phenomena
    - Automatic performance.
    - Contention scheduling without deliberate direction.
    - Deliberate conscious control.
  - Neuropsychological phenomena: The frontal lobes
  - Will
- Summary

Attention to Action: Willed and Automatic Control of Behavior

Donald A. Norman and Tim Shallice

Abstract

The major theme of the paper is that the primary role of attention is in the control of action. The basic idea is that human action sequences can run themselves off, efficiently, smoothly, without any need for deliberate attention. However, when modifications in a plan must be made, or when it is desired that some novel alternative action sequence be followed, or when it is desired to prevent some habitual act from occurring, then it is necessary for deliberate attentional intervention into the process.

We argue that most attentional conflicts occur with the initiation rather than the execution of actions. We suggest two levels of control: a contention scheduling mechanism that selects from among competing schemas; a supervisory attentional mechanism that biases the selection process. We propose that the supervisory attentional system is required where the action sequences are ill-learned or novel, where the action is highly critical or dangerous, or where planning is required. In other cases, selection is by contention scheduling alone. The result is three modes of the control of performance: automatic, contention scheduling without deliberate direction, and deliberate conscious control. Will becomes the application of attentional resources to the control of action.

Attention to Action: Willed and Automatic Control of Behavior*

Donald A. Norman
University of California, San Diego
La Jolla, California 92093

Tim Shallice
Medical Research Council
Applied Psychology Unit
Cambridge, England

During the performance of a complex action sequence, many different action components are likely to be active at any moment. This results from the fact that any particular action sequence is apt to be comprised of numerous components that are to be performed at different times. Moreover, in the conduct of their normal everyday activities, people often interweave a number of action sequences, doing several activities during overlapping time periods. Thus, an activity such as writing a letter can occupy considerable duration, and it is performed while engaged in other activities--listening to music, conversation, eating, or drinking. The writing of the letter itself has many different levels of operation, ranging from organizational aspects to the detailed motor movements that cause the appropriate marks to appear on the paper.

The initiation of any individual action component can be relatively straight-forward; do the action as soon as the appropriate triggering conditions occur. Complications occur in setting up the appropriate conditions and analyses. Complications also occur when numerous action components compete for overlapping use of some of some limited resource or for related structures, or when the processing structures and conditions are not set up sufficiently precisely that they will perform properly without some other level of monitoring and control. In this paper we propose a mechanism that allows for several control structures to interact in order to achieve smooth non-conflicting operation of the numerous action components that might be simultaneously awaiting their turn for action. We consider separately several different aspects of the situation: first, the nature of the knowledge structures that control actions; then what we call the "horizontal threads" that specify attentional control; and finally, a mechanism for conflict resolution.

---

*Research support to D.A. Norman was provided by the Office of Naval Research under contract N00014-79-C-0323. The collaboration was made possible by a grant from the Sloan Foundation to the Program in Cognitive Science at UCSD. Support was also provided by grant MH-15828 from the National Institute of Mental Health to the Center for Human Information Processing. We thank members of the "Skills" group of the Cognitive Science Laboratory at UCSD, especially David Rumelhart, Geoffrey Hinton, Wynne Lee, Jonathan Grudin, and Bernie Baars. We appreciate thoughtful reviews and comments by Roy D'Andrade, Steve Keele, John Long, George Mandler, and Peter McLeod. Approved for public release; distribution unlimited.

The Structure of the Paper

Our goal in this paper is to provide an account for both the experiential and the experimental phenomena of attention. We do this by examining what is known of the phenomena of attention and proposing a theoretical framework. The framework is structured around the notion of a set of active schemas, organized according to the particular action sequences of which they are a part, awaiting the appropriate set of conditions so that they can become selected to control action. The analysis is therefore centered around actions, primarily external actions, but the same principles apply to internal actions -- actions that involve only the cognitive processing mechanisms. The organization of this paper is first to specify the theoretical framework that will guide the later analysis, then to examine some of the phenomena of attention, first the experimental, then experiential, and finally, the neuropsychological. But before we start the theoretical framework, a brief review of some of the experiential phenomenology that surrounds the use of the term "automatic" is appropriate, for this conception plays a major role in the development of our ideas.

Automatic Performance

The term "automatic" is one of the more ubiquitous in the phenomenology of attention. However, the term has a number of different, though related, meanings. Experientially, there are at least three different meanings to the term. First, there is the way that certain tasks can be executed, without awareness of their performance (as in walking along a short stretch of flat safe ground). Second, actions may be both initiated and performed without deliberate attention or awareness (as in the automatic brushing away of an insect from one's arm). Third are cases like the orienting response, in which attention is drawn "automatically" to something, with no deliberate control over the direction of attention.

In addition, there are cases in which one can be passively aware of performing the actions, but without placing deliberate attention to them, and without any attempt to control them; an example of this latter type occurs in the performance of a skilled athletic task, where one might consciously be attending to the opponent, but be fully aware of the "automatic" hitting of the ball. Finally, within contemporary cognitive psychology, the term "automatic" is often defined operationally to refer to situations in which a task is performed without interfering with other tasks. In this situation, automatic is defined to mean that the task is performed without the need for limited processing resources (Shiffrin & Schneider, 1977).

The different uses of the term "automatic" require different explanations. We return to this point after we have outlined the theoretical framework. Basically, however, the model that we propose presumes that many action sequences are performed without any need for conscious awareness or attentional resources. It is only during the initiation or termination of sequences that attention is apt to be required, and then only with ill determined, or poorly learned tasks, or when the situation is determined to be critical or dangerous. The different senses of the term reflect different aspects of the mechanism, resulting from whether the control is entirely without attentional resources, or requires supervisory control, or attentional monitoring.

Relationship to Previous Work

The theoretical ideas developed in this model are consistent with a number of developments in the psychological literature on attention and the control of action. The emphasis that attentional limitations will have their major effect at the action end of analysis, with considerable parallel and non-conflicting processing prior to the initiation of action, is related to work of Keele (see the chapter by Keele & Neill, 1978). The basic notion that attentional processes play an overseeing role, activating whatever processing component is in need of supervisory assistance, has been suggested by LaBerge (1975), LaBerge and Samuels (1974); and Klein (1976). It is related to Posner's views of attentional biases providing costs and benefits in the production of responses (Posner, 1978). Our resource notions originate with Kahneman (1973), elaborated by Norman and Bobrow (1975) and Navon and Gopher (1979). Shallice's earlier work on the role of consciousness and action systems (Shallice, 1972, 1978), Norman (1981) on schemas and control structures, and Rumelhart and Norman (Note 3) on typing have played major roles in the theory that we have developed. The notion of schema has, of course, been around for a while, being introduced for motor actions by Bartlett (1932) and used for this purpose by Schmidt (1975). A more complete view of the views of schemas consistent with our usage is presented by Rumelhart and Ortony (1977).

A Theoretical Framework for the Analysis of Attention

Schemas and Processing Structures

Action sequences are complex ensembles of coordinated motor responses, oftentimes requiring some mental computation and decision making and considerable use of knowledge from the memory systems. We assume that specification of the components of actions and processing is done by means of numerous memory schemas, some organized into hierarchical or sequential patterns, others in heterarchical or independent parallel (but cooperating) patterns. Any given action sequence that has been well-learned is represented by an organized set of schemas, with one -- the source schema -- serving as the highest order control. The term "source" is chosen to indicate that the other component-schemas of an action sequence can be activated via the source. The procedural aspects of schemas require processing structures that carry out the operations specified, resulting in actions either upon an internal data base or upon the outside environment via the limbs and speech organs.

Conflicts in action sequences can arise for numerous reasons: several actions might require incompatible use of the same processing structures (as in simultaneously attempting to raise and to lower the hand); an action might require a difficult and unpracticed use of related structures (as in reaching for an object while making a precise movement of the leg); the action could require resources in excess of the capacity of a particular structure (as when attempting to do complex mental arithmetic while also retaining some items in short-term memory); or the result of one activity might preclude successful completion of another (as when eating dinner at location A precludes eating dinner at location B). Part of the difficulty in selection and scheduling of action components is to avoid incompatible or conflicting use of processing structures and to prevent the joint occurrence of other competitive activity. We propose that this occurs through selection, competition, and negotiation among schemas.

In the model there are three different states of a schema: dormant, activated, and selected. The state of dormancy is the normal, neutral state of the schema: a schema is dormant when it resides within the permanent memory structure, playing no role in the ongoing active processing of the moment. A schema is activated when it is set up, brought to a state of readiness and given an activation value. The activation value is determined by the combination of several factors, including the value given to it at set up by its source schema, the results of deliberate attentional activation or of motivation, the influence of the interaction with other activated schemas, and the goodness of match of the conditions within the trigger data base to the trigger conditions specified for the individual schema which determine the conditions under which it should be invoked. A schema is selected when its activation value is sufficiently high to exceed its own threshold. A selected schema controls actions, both internal processing and external movements of effectors.

We now describe these aspects in more detail and introduce an interaction of horizontal and vertical thread structures, and scheduling mechanisms.

Horizontal Threads

Start by considering a simple, self-contained, well-learned action sequence, perhaps the act of depressing a response switch upon the flashing of a particular light. This action sequence can be represented by a set of component schemas, triggered by the arrival of the appropriate perceptual event and resulting in the selection of the proper body, arm, hand, and finger movements to depress the button. Some or all of this processing sequence could be set up in advance by activation of the appropriate source schemas which in turn activates the detailed component schemas for carrying through the desired sequence of action upon the specified flash of light. [1] Whenever the action sequence is set up, its representation by means of action schemas constitutes a horizontal thread. The important point is that the processing structure can in principle be well specified.

---

[1] Just how much of the details of an action square can be preset is a point that needs to be empirically examined. The observation that the latency of a response is proportional to its complexity argues against

11

The general nature of the processing structure for a simple action sequence is shown in Figure 1. The essential components are shown in the horizontal grouping of component schemas for an action sequence: this is a horizontal thread of processing structures. In this example, four component schemas are shown, receiving information about sensory and motor activity from a "trigger data base" and making use of "psychological processing structures" in transforming their outputs into actions. For many actions, the specific processing units involved would be much more complex. Thus in the skilled operation of a motor skill such as writing to dictation, a complex set of specific processing units would be involved, including storage buffers of various sorts (see Ellis, 1980; Morton, 1980; Wing, 1978). Moreover in such skills, the conceptual relations among processing units and the initiation and execution of component schemas would be considerably more complex than the linear relation shown in Figure 1. However, we let the schematic conceptualization of the horizontal thread symbolize the specification of the processing structures, regardless of their actual complexity. A horizontal thread, therefore, stands for the specification of the components of the processing structures that control the over-learned aspects of action.

In general, there will be a number of different action sequences being performed at any given time, each specified by its own horizontal thread structure, as shown in Figure 2. The operations performed by the component schemas that comprise the horizontal threads include internal operations upon a memory data base, the formation and set up of other schemas or processing threads, and external operations such as speech or movement. Different component schemas might all access a common memory data base, and they might need to use the same processing mechanisms (e.g., the same memory structures or particular muscle groups). As a result, the different threads may interact with one another, as symbolized by the lines in Figure 2 interconnecting the threads. These interactions are important for the contention scheduling mechanism (to be described later).

Schema Selection Mechanisms

When numerous schemas are activated at the same time, some means needs to be provided for selection of a particular schema when it is required for its action sequence. At times, however, there will be conflicts among potentially relevant schemas, and so some sort of conflict resolution procedure must be provided. This is a common problem in any information processing system where, at any one moment, several

---

full, detailed specification of the motor schemas prior to the trigger signal (Kerr, 1975; Sternberg, Monsell, Knoll, & Wright, 1978). Our concern is independent of this consideration.

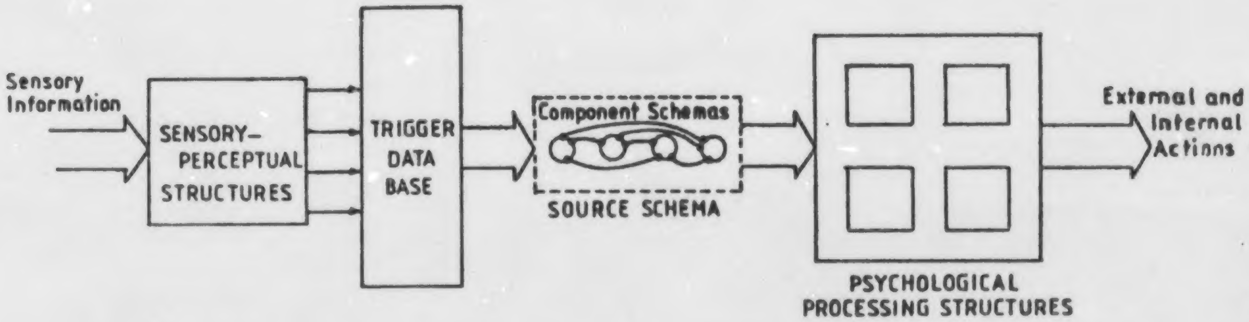
Sensory Information
SENSORY–PERCEPTUAL STRUCTURES
TRIGGER DATA BASE
Component Schemas
SOURCE SCHEMA
PSYCHOLOGICAL PROCESSING STRUCTURES
External and Internal Actions

Figure 1. A Horizontal Thread. For well-learned, habitual tasks an autonomous, self-sufficient strand of processing structures and procedures can usually carry out the required activities, without the need for conscious or attentional control. Selection of component schemas is determined, in part, by how well the "trigger conditions" of the schema match the contents of the "trigger data base." Such a sequence can often be characterized by a (relatively) linear flow of information among the various psychological processing structures and knowledge schemas involved: a horizontal thread.

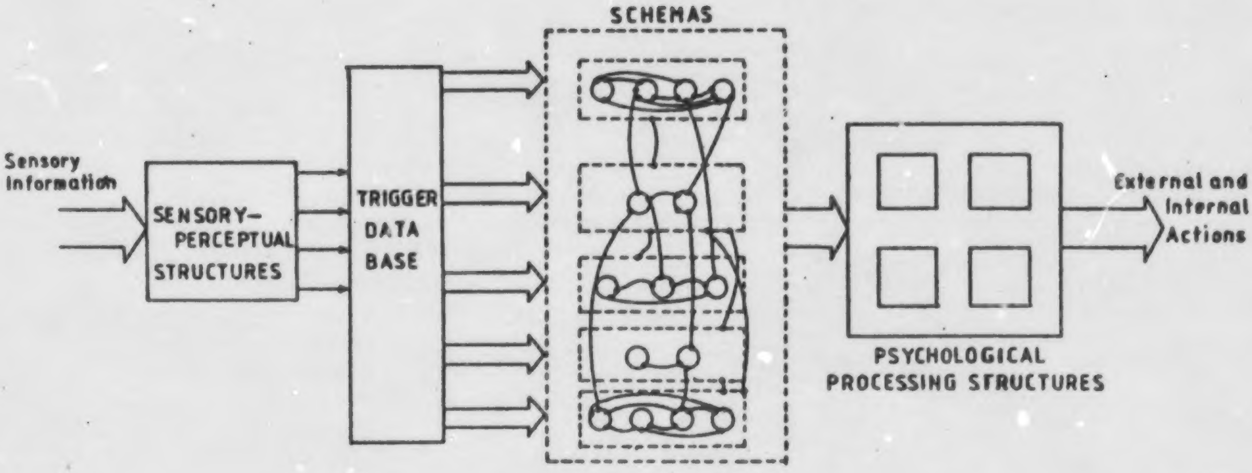
SCHEMAS
Sensory information
SENSORY-PERCEPTUAL STRUCTURES
TRIGGER DATA BASE
PSYCHOLOGICAL PROCESSING STRUCTURES
External and Internal Actions

Figure 2. Simultaneous Horizontal Threads. A person often performs several tasks at the same time, with the individual components of each task either being simultaneous or overlapping in time. Moreover, any given task may last for a considerable amount of time. This figure shows 5 different horizontal threads that might be active at one time. Some means of selecting the individual schemas at appropriate times while providing some form of conflict resolution becomes necessary. The interaction among the various horizontal threads needed for this purpose is indicted by the lines that interconnect schemas from different threads.

potential candidates for operation might require access to the same resources or might result in incompatible actions. (McDermott & Forgy, 1978, discuss this issue for production systems and Bellman, Note 1, discusses the problem with respect to animal behavior.)

The procedure we propose is constrained by the desire to transmit priorities by means of the single variable of amount of activation, a concept consistent with current psychological theory. We propose that the individual component schemas of the horizontal threads each have an activation value that is determined by a combination of factors, some that operate among schemas, some that result from special processes that operate upon the schemas.

We divide the activational influences upon a schema into three sets: vertical thread influences, contention scheduling influences, and trigger condition influences. Horizontal threads determine the organizational structure of the schemas and processing mechanisms for a particular action sequence. Vertical threads determine biases acting upon the selection process. Trigger conditions determine the appropriate timing for the initiation of schemas, and the contention scheduling mechanism combines these various influences and selects between candidate schemas where appropriate.

Vertical Threads

The horizontal thread specifies the organizational structure for the desired action sequence. However, a scheme may not be available that can achieve control of the desired behavior, especially when the task is novel or complex. In these cases, some additional control structure is required. The vertical thread influences provide one source of control upon the selection of schemas, operating entirely through the application of activation values to the schemas that can bias their selection by the contention scheduling mechanisms. There are two major vertical factors: motivational variables and attentional control. It is this latter factor that is the focus of this paper. The overall system is shown in Figure 3.

Attentional resources. Deliberate attentional control is the most important of the vertical thread influences. Here we postulate that a supervisory attentional mechanism is capable of monitoring the overall activity, then of supplying an increase or decrease in the activation values of the relevant schemas. Note that this is an indirect means of control of action. Attentional control is directed only at activation value, not directly at the selection. Moreover, it is control overlaid on top of the horizontal thread organization. When attentional activation of a schema ceases, the activational value will revert to its normal value.

Allport (1980) has criticized a wide range of attention theories for succumbing to what he calls the "GPLCCP" belief in a "General-Purpose Limited-Capacity Central Processor." We agree with much of his criticism, and our proposal is meant, in part, to overcome these

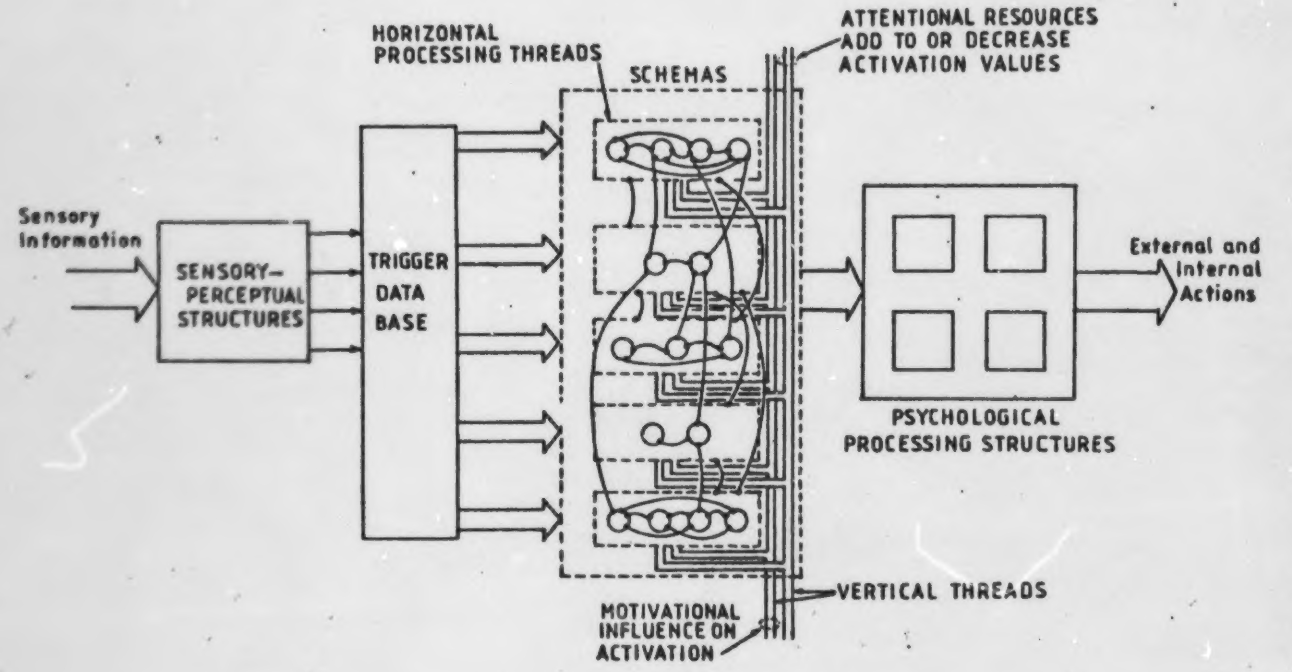
HORIZONTAL PROCESSING THREADS
SCHEMAS
ATTENTIONAL RESOURCES ADD TO OR DECREASE ACTIVATION VALUES
Sensory Information
SENSORY-PERCEPTUAL STRUCTURES
TRIGGER DATA BASE
PSYCHOLOGICAL PROCESSING STRUCTURES
External and Internal Actions
VERTICAL THREADS
MOTIVATIONAL INFLUENCE ON ACTIVATION

Figure 3. Vertical Threads. When attention to particular tasks is required, either because the components of the relevant horizontal threads are not sufficiently well specified or because some critical or dangerous situation is involved, then vertical thread activation comes into play. Attention operates upon schemas only through manipulation of activation values, increasing the values for desired schemas, decreasing (inhibiting) the values for undesired ones. Thus, attentional processes oversee and bias ongoing action by alteration of activation values. Motivational variables are assumed to play a similar role in the control of activation, but working over longer time periods.

criticisms. The horizontal processing threads represent particular strategies for performing tasks, making use of whatever processing structures seem needed. Several such threads may operate simultaneously provided that no related processing structures are simultaneously required and provided all the schemas involved are well-learned. Horizontal thread control is not subject to Central Processor limitations.

The model does contain a general purpose limited capacity mechanism, a supervisory attentional mechanism whose influence is felt by threading its way vertically across the active schemas. Mechanisms that are concerned with planning or monitoring of actions as contrasted with the detailed execution of the task solution can play important roles in overseeing the satisfactory operation of complex systems. Thus, such mechanisms have been incorporated into a number of models of problem solving programs (Boden, 1977; Fahlman, 1974; Sussman, 1975). In the present model, this mechanism only serves as a mediating influence; it can only modulate the flow of processing. Oftentimes, this modulation is critical for the successful operation, and whenever this is the case, central processing limitations can occur. But we presume the whole action system refines itself through experience, developing and adjusting the horizontal thread structures to minimize the need for central, vertical thread modulation.

Motivation. A second vertical thread component results from the effects of motivational factors. We take this to be a relatively slowly acting system, working primarily to bias the operation of the horizontal thread structures towards the long-term goals of the organism by activating source schemas (and through their selection, component schemas). Memory organizational procedures, for both storage and retrieval, are themselves horizontal thread structures, and so are susceptible to motivational biases.

Contention Scheduling

Simultaneously performed actions are sometimes in conflict with one another, but at other times can be jointly performed by cooperative action. In the case of conflict, when one task is started, something must prevent simultaneous performance of the other. But with cooperative tasks, the situation is quite different. Oftentimes the lower-level activities required for a single, higher level task are both cooperative and in conflict. To see this, consider how a skilled typist types the word "very" (using a standard typewriter keyboard): the positioning of the hands and fingers is both cooperative and competitive. The typing of the "v", "e", and "y" is cooperative; as the left index finger positions itself to type the "v", the middle finger of the left hand can start its movement towards the "e" and the right index finger can position itself for the "y". The hands and arm position themselves so as to assist in the finger movements. In contrast, the typing of the "v" and the "r" is competitive, both requiring conflicting use of the same finger. Analyses of high speed moving films of skilled typists indicates that both competitive and cooperative interactions occur (Gentner, Grudin, & Conway, Note 2).

To permit simultaneous action of cooperative acts and prevent simultaneous action of conflicting ones is a difficult job, for often the details of how the particular actions are performed determine whether or not they conflict with one another. Thus, whether or not several objects can be picked up at one time with the same hand is determined by the shape and location of the objects, the exact use that is made of the fingers, and the skill and experience of the person. Some sort of conflict resolution mechanism must come into play to resolve these issues.

We propose that the scheduling of actions takes place through what we call "contention scheduling," in which active schemas interact with one another through inhibition and, to a lesser extent, excitation of activation values. The result is competition for selection, modulated by the vertical thread activations, preventing competitive use of common or related structures and negotiating cooperative, shared use of common structures or operations where that is possible. (This principle is used by Rumelhart & Norman, Note 3, in a model of the hand and finger interactions in skilled typing.)

We assume that the initial activation values of component schemas are determined by means of their source schema. For example, when the source schema for a task such as driving an automobile has been selected, all its component schemas become activated, including schemas for such acts as steering, stopping, accelerating, slowing, overtaking, and turning. Each of these component schemas acts as a source schema, activating its component schemas (braking, changing gear, signalling). The one remaining consideration for the selection process is the satisfaction of the contention trigger conditions.

Trigger conditions. The determination of the activation level of a schema from higher-level systems is normally not sufficient to provide adequate timing of its selection. A particular action must be done at a time dependent upon the occurrence of appropriate environmental events. We propose that proper timing of schema selection depends upon the satisfaction of trigger conditions that specify the exact conditions under which selection is appropriate. Trigger conditions then contribute to the overall activation of the schema. How well the existing conditions match those trigger specifications determines the amount of activation contributed by this factor. (Attentional or motivational influences on activation value can override the effect of the trigger conditions, causing selection even in the absence of appropriate triggering conditions, or inhibiting selection even when there is perfect match of triggering conditions.)

The selection mechanism. There are two basic principles of the contention scheduling mechanism: first, the sets of potential source schemas compete with one another in the determination of their activation values; second, the selection takes place on the basis of activation value alone -- a schema is selected whenever its activation exceeds a threshold value. The threshold can be specific to the schema, and could become lower with use of the schema.

The competition is effected through lateral activation and inhibition among activated schemas. What degree of lateral inhibition exists between schemas on the model remains an open issue. Schemas which require the use of any common processing structures will clearly need to inhibit each other. Yet the degree of inhibition cannot be determined simply a priori. Thus some aspects of the standard refractory period phenomena can be plausibly attributed to such inhibition between schemas; explanations based upon conflicts in reponse selection fit the data well (Kahneman, 1973). However, one cannot just assume that responses by each of the two hands inevitably involve a common processing structure, as refractory period effects can disappear if highly compatible tasks are used (Greenwald & Shulman, 1973). On the model, as a task becomes better learned, the schemas controlling it could become more specialized in their use of processing structures, reducing potential structural interference and minimizing the need for mutual inhibition among schemas. At the same time, a factor that may operate to broaden lateral inhibitory interactions among schemas is the possibility of interactions among anatomically related subsystems, even when they are functionally distinct (Kinsbourne & Hicks, 1978).

Simultaneous selection of two schemas is unlikely for two reasons. First, if the two are at all incompatible, lateral inhibitory processes will tend to reduce the chance of at least one reaching threshold, as such effects magnify existing activation differences. The decrements will be greater the less well learned and, hence, the less specific the schemas are. Second, if deliberate attention is required to activate them sufficiently for selection, there will be competition for the limited resources of the supervisory attentional mechanism. This will be especially the case when tasks are not well learned. An ill-learned schema is itself a poorly integrated group of element schemas, and the supervisory attention mechanism will need to boost all the elements rather than the single whole. [2]

Note that the operation of a selected schema continues, unless actively switched off, regardless of what value its activation may have fallen to, until it has satisfied its goal, completed its operations, or until it is blocked when some resource or information is either lacking or is being utilized by some more highly activated schema. The activation value is important primarily in the selection process and when the selected schema must compete for shared resources, or in providing component schemas with initial activation values.

---

[2] Although simultaneous selection of two schemas is unlikely, simultaneous and synchronized performance of two action sequences is quite possible. If a person needs to perform two action sequences concurrently, both of which require conscious attention, then the simplest way to accomplish this is to establish a single, higher-order source schema that oversees both actions. Then, "attention" need be directed only at a single schema. If the underlying component schemas are automated (so they need not pass through contention scheduling), the result will be synchrony in action. Thus, in general, actions that require attentional control must either be performed with exact synchrony or in alternation, first one, then the other.

The scheduling is therefore quite simple and direct. No direct attentional control of selection is required (or allowed). Deliberate attention exerts itself indirectly through its effect on activation values. All the action, therefore, takes place in the determination of the activation values of the schemas.

Fine timing of operations. A critical component of the performance of many skills is fine timing of the operations. As the selection mechanism is now constituted, it cannot be counted upon to respond as precisely as is required. Even if triggering conditions were the same on different occasions, the contention selection mechanism would be apt to lead to variability in selection time, in part because it would be unreasonable to assume that other factors affecting activation level were constant.

One possibility is for the contention scheduling mechanism to be used only for initial selection of schemas and for crude timing. Precise timing would then be handled by means of specific triggering conditions at appropriate (low) levels of component schemas. Component schemas would then specify precise triggering conditions required for the actions under their control so that this level of control would not be done by the contention scheduling mechanism. This assumption also allows negotiations among schemas that are simultaneously operating to take place at this same level, so that one schema operates in such a way as to allow as much as possible of the other schema to be realized. Thus, if both a paper-picking-up schema and a pencil-picking-up schema are operative, the hand and finger configuration used for picking up the paper is likely to be modified so as to allow for the picking up of the pencil with the unoccupied fingers.

The Interpretation of the Phenomena

A key component of our theory is the role of attention to action. However, many of the experiments that have been performed to examine attentional phenomena have concentrated upon the role of attention in perception. As a result, many of these experiments do not test our ideas directly, but rather require explanations that cut across the several theoretical mechanisms that we describe.

One major class of phenomena do fit reasonably directly into the theoretical structure of the model. These are certain of the phenomenological aspects of attention, which are subject to reasonable agreement among observers, but available only through introspection rather then hard experimental evidence. In this section of the paper we discuss both the experimental and the introspective, phenomenological evidence.

The Experimental Literature

According to our framework for the control of action, simultaneous tasks interfere with one another in one of two ways: structural interference, when two horizontal threads overlap in their demands for processing structures; and attentional or resource interference, when attentional capacity is required from the vertical thread. One can attempt to control both these forms of interference, the structural interference through the appropriate design of the tasks that are to be performed at the same time, the attentional interference through sufficient (albeit lengthy and tedious) training on the tasks so that they can be performed "automatically," without the need for attentional biasing of the contention scheduling mechanism.

The best way to assess these ideas would be to have experiments that separate cleanly the demands placed upon horizontal and vertical thread components. We expect attentional limitations to effect only vertical thread processes and structural interference to affect only horizontal thread influences. However, as we discussed earlier, the structural interference between any pair of tasks cannot readily be determined a priori. Thus, the existence of interference between tasks does not speak directly to the theory. Rather, the more important predictions concern the conditions for which there will be no interference between tasks, as this requires that there be neither structural (horizontal) nor attentional (vertical) interference.

Simultaneous tasks. On the model, satisfactory performance of several simultaneous tasks depends upon lack of conflict of these tasks for any of several kinds of resources. The major prediction is that parallel dual-task performance should be most easily possible where only a single action sequence has to be initiated, that is, where only one schema has to be selected by contention scheduling. This fits with the results on monitoring (for review see Duncan, 1980). Thus Moray and Fitter (1973) and Sorkin and Pohlmann (1973) showed that monitoring for one of several possible signals (or signals over several possible channels) could be done satisfactorily, and that performance suffered only when simultaneous target detection and responding was required (in our terms, when there was initiation of actions).

Effective multiple channel monitoring depends on the activation of appropriate schemas by triggers that have been set up previously. There can, however, be difficulties in setting up the schema trigger conditions so as to match the expected signals unambiguously. For instance, how novel may the set of trigger conditions be? The Shiffrin and Schneider experiment (1977) in which the "automatic" attention was not possible when the target distractor relation was varied from trial to trial indicates that considerable practice is required to set up some trigger conditions. A second problem is that activation-based processes are liable to result in selection induced by stimuli similar to targets. The experiments by Treisman and Gelade (1980) seem to indicate that some situations which require the integration of "separable features" (where the background distractors may also contain these same features) present special difficulties of this sort. Thus, in some cases, effective target selection cannot be performed without internal processing actions. On the current model this is apt to require attentional activation of the processing schemas, thus forcing deliberate serial attention to be directed at the parts of the signal.

Simultaneous performance of tasks which require production of separate response streams is possible when two conditions are satisfied: first, the horizontal thread structures must be sufficiently developed that they can control the action sequences without deliberate supervision; second, there cannot be any structural interference. The first condition is obviously most apt to be satisfied when there has been considerable practice at the tasks. The second condition requires minimization of the overlap of use of common processing or control mechanisms. These conditions have been satisfied in a number of experiments that have examined performance of highly skilled, well-practiced people (Allport, Antonis, & Reynolds, 1972; McLeod, 1977; Spelke, Hirst, & Neisser, 1976).

Even in these situations, performance often deteriorates somewhat when a second task is added, although there appears to be no obvious grounds for structural or attentional interference. A common explanation of this finding is that there is an excess "overhead" associated with the performance of two tasks rather than one, an overhead that leads to a performance decrement regardless of the nature of the tasks (e.g., Allport, 1980). The model provides a possible source for this extra overhead: the supervisory attentional mechanisms. Unless the second task is extremely well-learned, there will be a need for extra source and component schemas that require vertical thread activation in their setup and in their selection. Thus, as Allport (1980) pointed out, in experiments involving piano playing conducted by Allport, Antonis, & Reynolds (1972), the one subject who showed no interference "was also the most competent of our pianists." The other subjects all found some technical challenge in the music such that "moments of emergency occurred," where recovery required some relatively unpracticed applications of keyboard technique and therefore, on our model, attentional resources.

Attentional demands at the initiation of actions. A major theme of the model is that attention is used primarily at the initiation of actions. This property of the model fits with the results of probe studies during the movement, where responses to probes at the start or end of movement are more delayed than those during execution (Posner & Keele, 1969; Ells, 1973). Indeed under certain conditions, although the start of the response shows the expected interferences, the large stages of response execution may not delay responses to probes at all (Posner & Keele, 1969). In interpreting these results it is important to realize that both the starting and stopping of a response requires selection by contention scheduling since, at one level, they involve different schemas. To stop a physical motor response requires that the limb motion be halted, which requires initiation of the action of muscle groups that can counter the momentum of the movement. It is just as hard to stop a movement as it is to start it. The exception is when movement is stopped by an external "stop." In this case, one would not expect attentional effects at the termination of the movement which is exactly what is found.

Moreover, attentional influences act only to bias the selection mechanisms, not to do the selection. One relevant study was performed by McLean and Shulman (1978) who found that when attention was first directed to the possibility of a signal and then distracted, there was a residual bias that remained from the initial investment of attention. We believe this finding to be more consistent with the view that attention can only bias processing actions than with the more usual view that attention selects processing.

A related prediction made by the model is that if the triggering potential of one type of stimulus is much stronger than that of a second type, then the former will be selected in contention scheduling even though the latter is being deliberately attended to. In this situation, triggering activation is more powerful than activation from the supervisory attentional mechanisms. One set of such findings comes from the literature on selective attention in which an attempt is made to keep the subject concentrating upon a primary task while other signals are presented. A classic example of the difficulty of doing this is the Stroop phenomenon. Certain classes of words presented upon a secondary channel can intrude upon or bias primary task performance, such as a word that fits within the context of the primary channel, or that has been conditioned to electric shock, or that has high emotional value (such as one's own name). Performance of the other task is impaired when the interrupt occurs. In terms of our model, these "intrusions" result from data-driven entry of action schemas into the contention scheduling mechanism and their selection there due to the strongly activating properties of such triggers. These intrusions, therefore, are similar to the form of action slip found in "capture errors," to be discussed later.

Experiential Phenomena

An important aspect of attention and the control of action is its phenomenology. The role of the conscious awareness of directed, controlled attention to events and to actions, and the nature of those situations in which one can act without awareness or deliberate intention to act, are all of direct relevance to the theoretical structures described here. Indeed, the theory was developed with explanation of the experiential aspects of attention as much of a goal as the results of controlled experimentation.

Phenomenal reports of action provide evidence for qualitative distinctions among different types of experience. We do not discuss the correspondence between consciousness and information processing mechanisms and functions (but see Shallice, 1972, 1978). To make differentiations between types of experience, it is sufficient to assume that the supervisory attentional mechanism has the possibility for access to information about the schemas selected in contention scheduling, and that reflection on action involves this process.

Automatic performance. In an earlier section we identified several different aspects of the term "attention." On the model, there are correspondences for all these related aspects of the term. Some actions, such as the decision to walk across a certain path, might be initiated with directed attention and clear awareness, but the performance need not depend upon the supervisory attentional mechanism nor interfere with any other activated schema in contention scheduling; hence, the automatic nature of the performance. In other actions, such as the brushing away of an insect, even the initiating of the act (through data-driven excitation of a source schema) might not depend upon the supervisory attentional mechanism. In this case, there could be a complete lack of awareness of both the initiation and the action. These cases also correspond to the operational definition, for there is no demand upon the attentional resources.

There are, however, cases in which one experiential sense of "automatic" does not correspond to "automatic" in the operational sense. Thus, the orienting response is phenomenologically automatic, as there is no deliberate attentional control over schema activation and selection. However, schema selection may very well interfere in contention scheduling with the operation of tasks being performed at the same time; in these cases, the operational criterion for automaticity is not met.

A specially interesting case of automaticity is where at one time the action did require conscious direction for its performance, but now no longer does so. As William James (1890) pointed out at some length, the common denominator of such actions is that they are habitual, frequently performed, usually in a relatively fixed format. Why should the ability to perform such tasks automatically only occur when they are well learned? On the model, this is because newly learned actions are apt to be ill-specified. Their schemas are relatively small, encompassing relatively specialized sub-actions. Moreover, their triggering conditions are apt to be ill-specified, not well matched to the actual conditions that occur. As a result, continual monitoring is required by the attentional mechanisms, and selection must often be forced (or delayed) by the application of deliberate attentional activation. Well-learned actions are apt to be well specified, with their schemas encompassing large, organized units of behavior, and with their triggering conditions well-matched to the situation. As a result, once their schemas have been selected, they can maintain control effectively for longer periods.

The amount of interference that a task presents to other tasks performed at the same time will also decrease as learning improves. This decrease comes about for three different reasons. First, as the schemas become better specified, there will be fewer gaps and fewer weaknesses that need the supervisory processes for proper control. Second, as the triggering conditions become better matched to the situation, they are less likely to need supervisory attentional resources. Third, general schemas are apt to make more general demands on processing structures, making it more likely that there will be conflicts with other action sequences. As the actions become specified more precisely, they involve a smaller fraction of the psychological processing structures. This too reduces the conflicts in performance, by decreasing the possibility of lateral inhibitory conflicts during contention scheduling.

Contention scheduling without deliberate direction. It is possible to be aware of performing an action without paying active, directed attention to it. This corresponds to situations in which the selection of a schema is accomplished by contention scheduling without the involvement of the supervisory attentional mechanism. The most general situation of this type is in the initiation of routine actions. Phenomenally, this corresponds to the state that Ach (1905) describes as occurring after practice in reaction time tasks. Over the first few trials, he said, the response is preceded by awareness that the action should be made, but later there is no such awareness, except if preparation has been inadequate. By then, the stimulus triggers the appropriate schema without the involvement of the supervisory attentional mechanisms. In well-learned tasks, the subject experiences the response as proceeding with "an awareness of determination" even if it is not immediately preceded by any experience of intention to act. In our terms, this is because the schema controlling action had previously been activated using the supervisory attentional mechanisms. but then the schema selection occurs automatically when the proper stimulus conditions occur.

Whenever contention scheduling takes place without any present or prior involvement of the supervisory attentional mechanisms, even the "awareness of determination" is absent. Data-driven triggers act in this way. Because there is no involvement of the supervisory attentional mechanism, and hence no monitoring, errors can easily occur. Two such errors are the form labelled "data-driven errors" and "capture errors" (Reason, 1979; Norman, 1981). Both classes of errors occur when the schema that was intended to control action is replaced by another one, leading to an unintentional result. In the case of "data-driven" errors, the unintended schema is activated by perceptual information (by newly arriving sensory data). In the case of "capture errors," the unintended schema shares considerable features with the intended one, and in addition, is the more frequently performed action sequence. In either case, one may find oneself doing a totally unexpected set of actions, much to one's own dismay.

Take for example, Reason's description of the person who went to the garage to drive to work and found that he had "stopped to put on my Wellington boots and gardening jacket as if to work in the garden" (Reason, 1979). Or take the person described by William James who went to the bedroom to change for dinner and ended up undressed, ready for bed. Students of abnormal behavior may wish to point out that there is seldom a single explanation for behavior, that there are many interacting causes. Thus, the person who found himself gardening might also have (subconsciously) wished to avoid going to work. So too with the person described by James; the dinner may have been unwelcome. Our contention scheduling system is deliberately designed with these issues in mind. We postulate that selection results from the combination of numerous factors. Thus, it is quite possible that the sight of the gardening boots or the act of undressing (to change one's clothes) would not by themselves have been sufficient to have selected the discrepant behavior. Similarly, the hidden wishes, whether conscious or not, would not by themselves have been sufficient to cause the behavior. But the fortuitous combination of the wishes and the situation were sufficient to cause selection of the schemas.

Deliberate conscious control. A critical separation on the model is between action initiated through contention scheduling without the involvement of the supervisory attentional mechanism, and action initiated with the involvement of this mechanism. This distinction corresponds closely to William James's (1890) distinction between "ideo-motor" and "willed" acts. To James, "wherever movement follows *unhesitatingly and immediately* the notion of it in the mind, we have ideo-motor action. We are then aware of nothing between the conception and the execution." These "ideo-motor" actions (a category which does not exclude "awareness of determination") corresponds directly to our idea of contention scheduling without conscious direction. His concept of actions involving "an additional conscious element in the shape of a fiat, mandate, or expressed consent" corresponds to cases where we believe the supervisory attentional mechanism to be operative.

Experientially, a number of different sorts of tasks appear to require a considerable amount of deliberate attentional resources. These tasks fit within the following categories:

(a) they involve planning or decision-making,

(b) they involve components of trouble shooting,

(c) they are ill-learned or contain novel sequences of actions,

(d) they are judged to be dangerous or technically difficult,

(e) they require overcoming a strong habitual response or resisting temptation.

The general principle involved is that these are special situations in which the uncontrolled application of a horizontal processing thread through the contention scheduling mechanism is apt to lead to error. The supervisory attentional mechanisms allow more control over the sequence of actions to be performed than is possible through horizontal thread direction alone.

Vertical thread influences are necessary because individual schemas are limited in what they can foresee and control. Further, although the contention scheduling mechanism allows selection among schemas, it does not allow for integration of information across them. Given the variety and power of human capacities, it would appear that mechanisms must exist which have available to them information about the varied needs and capacities of the organism, including source schema, the ability to monitor the operation of schemas, and the power to initiate the construction of new schemas out of existing ones.

Planning and decision making are processes that operate in the formation of intentions that are not routine. In our terms, we plan or decide when it is clear that no existing schemas are sufficient to satisfy a particular goal. In these cases, information must be from more than one schema, or new schemas must be formed. This requires involvement of the supervisory attentional mechanisms. We assume though, that these general powers are bought at the cost of speed. Use of the supervisory attentional mechanisms then provides both benefits and costs. The benefits derive from the increased processing power brought to bear on the problem at hand; in general, judgments which it controls will be superior to the unguided selections of contention scheduling. The costs result from the slowness, leading to difficulty in the control of rapid, skilled actions, and to seriality in the control of what would otherwise be parallel acts.

We define trouble shooting to be the application of planning and decision-making processes to actions already in progress. It occurs when an unexpected error occurs in the operation of an action (see Mandler, 1975). When a particular, specialized component schema has failed, one solution is to replace it with a more general one. More general schemas are apt to require selection through contention scheduling and vertical thread control by the supervisory attentional mechanisms.

The performance of ill-learned or novel skills requires what Fitts and Posner (1967) called "the early or cognitive phase ... in which it is necessary to attend to cues, events, and responses that later go unnoticed" (pp.11-12). These are situations which require attentional control because neither appropriate schemas nor their triggers have been developed. With dangerous and technically difficult situations, error is relatively costly. In this situation, one wishes to guard against the vagaries of contention scheduling and enforce selection of the most appropriate schema by means of strong activation. Similarly, in overcoming habitual responses or in resisting temptation, the appropriate schema must be strongly activated and the others strongly inhibited, else a more usual, but inappropriate act, might be selected by the normal operation of contention scheduling.

There can be costs in using deliberate control in a task that is normally performed automatically. Activation of schemas in contention scheduling by the supervisory attentional mechanism has the effect of reducing the influence of activation from triggering stimuli and other schemas in the selection process: subtle environmental control is lost. In addition, action execution must proceed unit by unit, each schema awaiting its turn for receiving attentional biases. The result is a lack of smoothness and a slowing of performance.

Neuropsychological Phenomena: The Frontal Lobes

There are strong correspondences between functions of our supervisory attentional mechanism and those ascribed by Luria (1966) to the prefrontal regions of the brain. If the supervisory attentional mechanism were damaged, the resulting behavior would be similar to the behavior of patients with lesions to the prefrontal regions.

A deficit in planning corresponds to Luria's clinical characterization of the frontal lobe syndrome, which has been supported in a number of experimental studies involving maze learning, complex visual-constructive tasks, and complex arithmetical problem-solving (see Walsh, 1978). The simplest example of a planning disorder is the finding of Gadzhiev (see Luria, 1966) that frontal patients when presented with a problem tend to miss out the initial assessment of the situation. Frontal lobe patients have also been characterized clinically as having deficits in initiative, in dealing with novelty, and of judgment (Penfield & Evans, 1935; Goldstein, 1936).

Patients with frontal lobe lesions have difficulties with error correction. The Wisconsin card-sorting test involves multi-dimensional stimuli where the patient must switch from sorting according to one dimension to sorting according to another. In this task frontal patients show a strong tendency to perseverate in sorting on the previously correct dimension, even when they are told they are wrong (Milner, 1964; Nelson, 1976). Perret (1974) found that patients with frontal lobe lesions are the most impaired group on the Stroop test. This is a task in which the usual response to a stimulus is not the desired one -- habitual responses must be suppressed. In this situation deliberate attentional control is required, but this in general presents especial difficulty for frontal lobe patients.

The failure to overcome an habitual response tendency is one side of the general effect that should occur on our model if the supervisory attentional mechanisms are damaged. In this case, behavior will be left under the control only of the horizontal thread structures, plus contention scheduling. In the examples above where one schema is more strongly activated than the others, it will be difficult to prevent it from controlling behavior. By contrast, when several schemas have similar activation values one should obtain another clinical characteristic of frontal patients: an instability of attention and heightened distractability (see Walsh, 1978). This apparent contradiction between increased perseveration and increased distractability results from failure of a single mechanism. Both results are observed in animals with prefrontal lesions (see Brush, Mishkin, & Rosvold, 1961; Pribram, 1973).

If the properties of the supervisory attentional mechanism seem to correspond fairly well with neuropsychological evidence, does the same apply to the properties of contention scheduling? One possible relation is between the lateral inhibitory and threshold properties of contention scheduling and certain properties of the basal ganglia, thought because of their role in the aetiology of Parkinson's disease to be involved in the initiation of action (see also Stein, 1978). Moreover the basal ganglia are innervated by dopamine systems, which it has recently been claimed mediate the selection of behaviors through a lateral inhibitory mechanism somewhat analogous to contention scheduling (see Joseph, Frith, & Waddington, 1979) and which when they malfunction (as in amphetamine psychosis) lead to disorders which could well be at the level of the selection of action (see Lynn & Robbins, 1975).

Will

We propose that "will" be the direction of action by direct conscious control through the supervisory attentional mechanism. This definition is consistent both with the popular meaning of the term and with the discussions of will in the earlier psychological literature. Thus, strongly resisting a habitual or tempting action or strongly forcing performance of an action that one is loathe to perform seems to be prototypical examples of the application of will. The former would appear to result from deliberate attentional inhibition of an action schema, the latter from deliberate activation. James (1890) drew the contrast between "what happens in deliberate action" where will is involved and actions that do not require will, where the responses followed "unhesitatingly and immediately the notion of it in the mind" (ideo-motor actions). Situations in which there is no need for will are those where there "seems to be the absence of any conflicting idea."

In our view, will varies along a quantitative dimension corresponding to the amount of activation or inhibition required from the supervisory attentional mechanisms. The assumption that this activation value lies on a continuum explains why the distinction between willed and ideo-motor actions seems quite clear when considering extreme actions, but becomes blurred when considering those that require very little attentional effort. Thus, introspection fails in determining whether or not will is involved in the voluntary lifting of the arm. But there is no need to make a distinction if this act is simply identified as being near the zero point of the quantitative scale of attentional activation.

The idea that will corresponds to the output of the supervisory attentional mechanisms has certain other useful consequences. Consider the errors that occur with brief lapses of attention, when there is a failure to sustain will adequately. One type of error results following a decision not to do a step within a habitual sequence of actions. To eliminate the step requires deliberate (willful) inhibition of the relevant schema. If there is a momentary lapse of attention to the deliberate inhibition, the step may get done anyway. Closely related is the error that occurred to one of us. Having decided not to take another bite of a delicious, but extremely rich dessert, with only a brief lapse of attention, the cake got eaten.

Certain aspects of will require elaboration of our approach. In some circumstances an action may seem to require no will at all, yet at other times, require extreme demands. Thus, getting out of bed in the morning is at times an automatic act, at other times requires great exertion of will. One explanation for this observation is that activation of an action schema by the attentional mechanisms necessarily involves knowledge of consequences. When these are negative, they lead to inhibition of the source schemas which then must be overcome. In some cases, the self-inhibition can be so intense as to prevent or at least make very difficult the intended act. Thus, inflicting deliberate injury to oneself (as in pricking one's own finger in order to draw blood) is a difficult act for many people.

The elicitation of strong activation from the supervisory attentional mechanism is not necessarily unpleasant. Indeed, many sports and games seem to be attractive because they do necessitate such strong activation. In this case "concentration" is perhaps the more appropriate experiential equivalent rather than "will." In addition, will is not just a matter of attention to actions. As Roy D'Andrade (personal communication) has pointed out, a willed act demands not only strong attentional activation, it also depends on the existence of a "mandated decision," independent of one's attending, a conscious knowledge that the particular end is to be attained. This mandate, in our view, would be required before the supervisory attentional mechanisms will produce their desired activation output. However the critical point for the present argument is that the phenomenal distinction between willed and ideo-motor acts flow from separation of the supervisory attentional mechanisms from the systems they oversee. The phenomenology of attention can be understood through a theory of mechanism.

30

Summary

We present a possible framework for considering the role of attention in the control of action. In this, we have emphasized several things. First, because people usually do numerous activities during a given time period, a major concern for the control of action becomes how the selection of the individual components occurs at appropriate times, allowing co-operative actions to co-occur and avoiding conflicting ones. Second, there is the importance of the initiation of action. We assign the basic role of the attentional mechanisms to the initiation of action (as opposed to perceptual analysis -- where the bulk of the experiments have been performed -- or to thought and decision processes). By "action" we include the initiation of both internal processing actions and external control of effectors. Third, we emphasize that many activities can be carried out autonomously, without the need for conscious or attentional control, by means of well-specified, horizontal thread processing components. It is only in cases where the action sequences are ill-specified, or in situations that are judged to be critical or dangerous that deliberate attentional control is required. In this case, we suggest that supervisory attentional mechanisms of limited capacity oversee the operation of the system, monitoring for the success of the activity, and biasing the selection and suppression of component schemas by altering the activation values of those schemas. We specify that such attentional control does not act directly, but only indirectly through the mediation of activation value.

By this scheme, there are two forms of interference likely to be encountered in the production of simultaneous tasks. The two forms correspond to (a) interference among horizontal threads when they must compete for use of overlapping processing mechanisms and to (b) interference among the vertical thread activations when they must be produced by the supervisory attentional mechanism. The first form of interference -- horizontal thread interference -- is related to "structural interference," but on our approach this is mediated by the operation of a mechanism, contention scheduling, that selects from potentially competing actions using only the activation levels of the schemas that control them. The second form of interference -- vertical thread interference -- is more a form of "resource interference."

There are two different modes for the control of action and, like the distinction between forms of interference, they also correspond to the difference between horizontal and vertical thread control. Thus, when processing sequences are sufficiently well specified that they can be controlled entirely by horizontal thread operations, they correspond to "automatic" actions. When conditions do not permit unsupervised horizontal control (or when the person deliberately invokes attentional processes to the action sequence), then the operations correspond to processing under "conscious control" or "willed" action.

Reference Notes

1. Bellman, K. *The conflict behavior of the lizard*, Sceloporus Occidentalis, *and its implication for the organization of motor behavior*. Unpublished doctoral dissertation. University of California, San Diego, 1979.

2. Gentner, D. R., Grudin, J., & Conway, E. Finger movements in transcription typing. Center for Human Information Processing, University of California, San Diego. CHIP Technical Report, 1980.

3. Rumelhart, D. E., & Norman, D. A. Simulating a skilled typist: A study of skilled cognitive-motor performance. Unpublished manuscript, The University of California, San Diego, 1980.

References

Ach, N. *Uber die Willenstatigheit und das Denken*. Gottingen: Vardenhoek, 1905.

Allport, D.A. Attention and performance. In G.L. Claxton (Ed.), *New directions in cognitive psychology*, London: Routledge, 1980.

Allport, D. A., Antonis, B., & Reynolds, P. On the division of attention: A disproof of the single channel hypothesis. *Quarterly Journal of Experimental Psychology*, 1972, *24*, 225-235.

Bartlett, F.C. *Remembering: An experimental and social study*. Cambridge: Cambridge University Press, 1932

Boden, M. *Artificial intelligence and natural man*. New York: Basic Books, 1977.

Brush, E. S., Mishkin, M., & Rosvold, H. E. Effects of object preferences and aversions on discrimination learning in monkeys with frontal lesions. *Journal of Comparative and Physiological Psychology*, 1961, *54*, 319-325.

Duncan, J. The locus of interference in perception of simultaneous stimuli. *Psychological Review*, 1980, *87*, 272-300.

Ellis, A. W. Slips of the pen. *Visible Language*, 1980, *3*, 265-282.

Ells, J. G. Analysis of temporal and attentional aspects of movement control. *Journal of Experimental Psychology*, 1973, *99*, 10-21.

Fahlman, S. E. A planning system for robot construction tasks. *Artificial Intelligence*, 1974, *5*, 1-50.

Fitts, P. M., & Posner, M. I. *Human performance*. Belmont, CA.: Brooks-Cole, 1967.

Goldstein, K. The Significance of the frontal lobes for mental performance. *Journal of Neurology and Psychopathology*, 1936, *17*, 27-40.

Greenwald, A.G., & Shulman, A.G. On doing two things at once. II. Elimination of the psychological refractory period. *Journal of Experimental Psychology*, 1973, *101*, 70-76.

James, W. *The principles of psychology*. New York: Holt, 1890.

Joseph, M.H., Frith, C.D., & Waddington, J. L. Dopaminergic mechanisms and cognitive deficits in schizophrenia. *Psychopharmacology*, 1979, *63*, 273-280.

Kahneman, D. *Attention and effort*. Englewood Cliffs, N.J.: Prentice-Hall, 1973.

Kelso, J. A. S., Southard, D. L., & Goodman, D. On the coordination of two-handed movements. *Journal of Experimental Psychology: Human Perception and Performance*, 1979, *5*, 229-238.

Kerr, B. Processing demands during movement. *Journal of Motor Behavior*, 1975, *7*, 15-27.

Kinsbourne, M., & Hicks, R. E. Functional cerebral space: A model of overflow, transfer and interference effects in human performance. In J. Requin (Ed.), *Attention and performance Vol. 7*. Hillsdale, N.J.: Erlbaum, 1978

Klein R.M. Attention and movement. In G. Stelmach (Ed.), *Motor Control: Issues and Trends*, New York: Academic Press, 1976.

LaBerge, D., & Samuels, S.J. Toward a theory of automatic information processing in reading. *Cognitive Psychology* 1974, *6*, 293-323.

LaBerge, D. Acquisition of automatic processing in perceptual and associative learning. In P.M.A. Rabbitt & S. Dornic (Eds.), *Attention and performance*, *Vol. 5*. London: Academic Press, 1975.

Luria, A. R. *Higher cortical functions in man*. London: Tavistock, 1966.

Lynn, M., & Robbins, T. The action of central nervous system drugs: A general theory concerning amphetamine effects. In W.B. Essmann & L. Valzelli (Eds.), *Current developments in psychopharmachology*, *Vol. 2*. New York: Spectrum, 1975.

Mandler, G. *Mind and emotion*. New York: Wiley, 1975.

McDermott, J., & Forgy, C. Production system conflict resolution strategies. In D.A. Waterman & F. Hayes-Roth (Eds.), *Pattern-directed inference systems*. New York: Academic Press, 1978.

McLean, J. P., & Shulman, G. L., On the construction and maintenance of expectancies. *Quarterly Journal of Experimental Psychology*, 1978, *30*, 441 - 454.

McLeod, P. D. A dual task response modality effect: Support for multiprocessor models of attention. *Quarterly Journal of Experimental Psychology*, 1977, *29*, 651 - 658.

Milner, B. Some effects of frontal lobectomy in man. In J.M. Warren & K. Akert (Eds.), *The frontal granular cortex and behaviour*. New York: McGraw Hill, 1964.

Moray, N., & Fitter, M. A theory and the measurement of attention. In S. Kornblum (Ed.), Attention and performance, IV. New York: Academic Press, 1973

Morton, J. The logogen model and orthographic structure. In U. Frith (Ed.), Cognitive approaches in spelling. London: Academic Press, 1980.

Navon, D., & Gopher, D. On the economy of the human processing system: A model of multiple capacity. Psychological Review, 1979, 86, 214-255.

Nelson, H. A modified card sorting test sensitive to frontal lobe defects. Cortex, 1976, 12, 313-324.

Norman, D. A., Categorization of action slips. Psychological Review, 1981, in press.

Norman, D. A., & Bobrow, D. G. On data-limited and resource-limited processes. Cognitive Psychology, 1975, 7, 44-64.

Penfield, W. & Evans, J. The frontal lobe in man: A clinical study of maximum removal. Brain, 1935, 58, 115-133.

Perret, E. The left frontal lobe of man and the suppression of habitual responses in verbal categorical behaviour. Neuropsychologia, 1974, 12, 323-330.

Posner, M.I., & Keele, S.W. Attention demands of movement. In Proceedings of the 16th International Congress of Applies Psychology. Amsterdam: Swets and Zeitlinger, 1969.

Posner, M. I., Chronometric explorations of mind. Hillsdale, N. J.: Erlbaum, 1978

Pribram, K.H. The primate frontal cortex-executive of the brain. In K.H. Pribram & A.R. Luria (Eds.), Psychophysiology of the frontal lobes. New York: Academic Press, 1973.

Reason, J. T. Actions not as planned. In G. Underwood & R. Stevens (Eds.), Aspects of consciousness. London: Academic Press, 1979.

Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Erlbaum Associates, 1977.

Schmidt, R.A. A schema theory of discrete motor skill learning. Psychological Review, 1975, 82, 225-260.

Shallice, T. Dual functions of consciousness. Psychological Review, 1972, 79, 383-393.

Shallice, T., The dominant action system: An information-processing approach to consciousness. In K. Pope & J. E. Singer (Eds.), The flow of conscious experience. New York: Plenum, 1978

Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. Psychological Review, 1977, 84, 127-190.

Sorkin, R. D., & Pohlmann, L. D. Some models of observer behavior in two-channel auditory signal detection. Perception and Psychophysics, 1973, 14, 101-109.

Spelke, E., Hirst, W., & Neisser, U. Skills of divided attention. Cognition, 1976, 4, 205-230.

Stein, J. The long loop motor control in monkeys. In J.E. Desmedt (Ed.), Progress in Clinical Neurophysiology Vol. 4. Basel: Karger, 1978.

Sternberg, S., Monsell, S., Knoll, R.L., & Wright, C.E. The latency and duration of rapid movement sequences. In G.E. Stelmach (Ed.), Information processing in motor control and learning. New York: Academic Press, 1978.

Sussman, G. J., A computational model of skill acquisition. New York: American Elsevier, 1975.

Treisman, A. M., & Gelade, G. A feature-integration theory of attention. Cognitive Psychology, 1980, 12, 97-136.

Walsh, K.W. Neuropsychology: A clinical approach. Edinburgh: Churchill Livingstone, 1978.

Wing, A.M. Response timing in handwriting. In G.E. Stelmach (Ed.), Information processing in motor control and learning. New York: Academic Press, 1978.

36

MULTIPLE RESOURCES IN TIME-SHARING[1]

Chris Wickens
University of Illinois

Recent research has addressed the question of whether human attention, or information processing resources can be effectively modelled as a multidimensional commodity (e.g., Navon & Gopher, 1979; Wickens, 1980). The heuristic model that we have proposed, as a framework for hypothesis formulation and test of the multiple resource concept suggests that processing resources are defined by processing stages at two levels (perceptual/central processing versus response processes), processing codes (verbal versus spatial), and processing modalities (auditory versus visual input, vocal versus manual response).

If validated, this configuration has both theoretical implications and practical implications for systems design and assessment. These will be addressed in turn, followed by a discussion of validating experimental data collected in our laboratories.

Theoretical Implications

If separate resources do underlie processing, then two implications follow directly: (1) Two tasks will be more efficiently time-shared to the extent that they demand non-overlapping resource "pools." Whether perfect time-sharing will result if the two pools are completely disjunctive, depends upon the demand (task difficulty) of the component tasks. (2) When the difficulty of one member of a time-shared pair is increased, its effect upon performance of the concurrent task will be a function of two factors: (a) the extent to which the subject treats the manipulated task as "primary" (tries to maintain constant performance despite the demand increase); (b) the extent to which the resources consumed by the demand increase are also those deployed in performance of the concurrent task. (Note that the effects of difficulty may be used interchangeably with the effects of automation or practice, since this variable may also have its influence on separate capacity pools. That is, automation of perceptual processing is functionally equivalent to a decrease in perceptual difficulty. Development of a motor program is equivalent to a decrease in response difficulty.)

---

[1] Contractual support for the research described in this report was provided by the Office of Naval Research (Gerald Molecki, Technical Monitor), Air Force Office of Scientific Research (Dr. Alfred Fregby, Technical Monitor), and DARPA (Dr. Craig Fields, Technical Monitor). Professor Emanuel Donchin was a collaborator in all aspects of the Event Related Potential Research. Richard Gill and Major William Derrick (USAF) have also assisted in many phases of the research reported.

While the two assertions above represent a skeletal statement of the elements of multiple resource theory, it is necessary also to consider in more detail the defining properties of a "resource" (in a time-sharing context), and to emphasize the distinctions that can be drawn between a multiple resource viewpoint, and the conception of a single resource with structural interference (Kahneman, 1973). A resource here refers to any commodity that can be divided or shared between tasks, and whose division can be modulated in continuous fashion, according to an allocation policy. In contrast, a dedicated processing structure must service only one task at a time, and therefore behaves in a two-state all-or-none fashion. (This contrast leaves ambiguous, the designation of a multiplexing system that can rapidly switch between two tasks, and whose dwell time can be adjusted in continuous analog fashion. Obviously, at a high enough switching frequency the dedicated processing mechanism becomes empirically indistinguishable from the resource). At issue then is whether there exists only one commodity with resource-like properties, or a number of such commodities. The POC (performance operating characteristic) methodology (Norman & Bobrow, 1975; Navon & Gopher, 1979) has provided a convenient framework for comparing these hypotheses.

If a multiple resource framework is adopted (as I have done here), it is necessary to specify the composition, or functionally defining properties of the resource pools, and the interrelationships between them. Assuming as I have done that resources may be defined by stages, codes, and modalities (Figure 1), (Wickens, 1980), three alternative conceptions are possible.
(1) An independence conception argues that the resources within each cell of the matrix in Figure 1 are disjoint and independent from the resources in all other cells. This assumption seems to be clearly untenable, as for example numerous experimental findings support the conclusion that auditory and visual processes will compete, as will spatial and verbal (albeit to a lesser extent than auditory-auditory, visual-visual, or spatial-spatial and verbal-verbal).

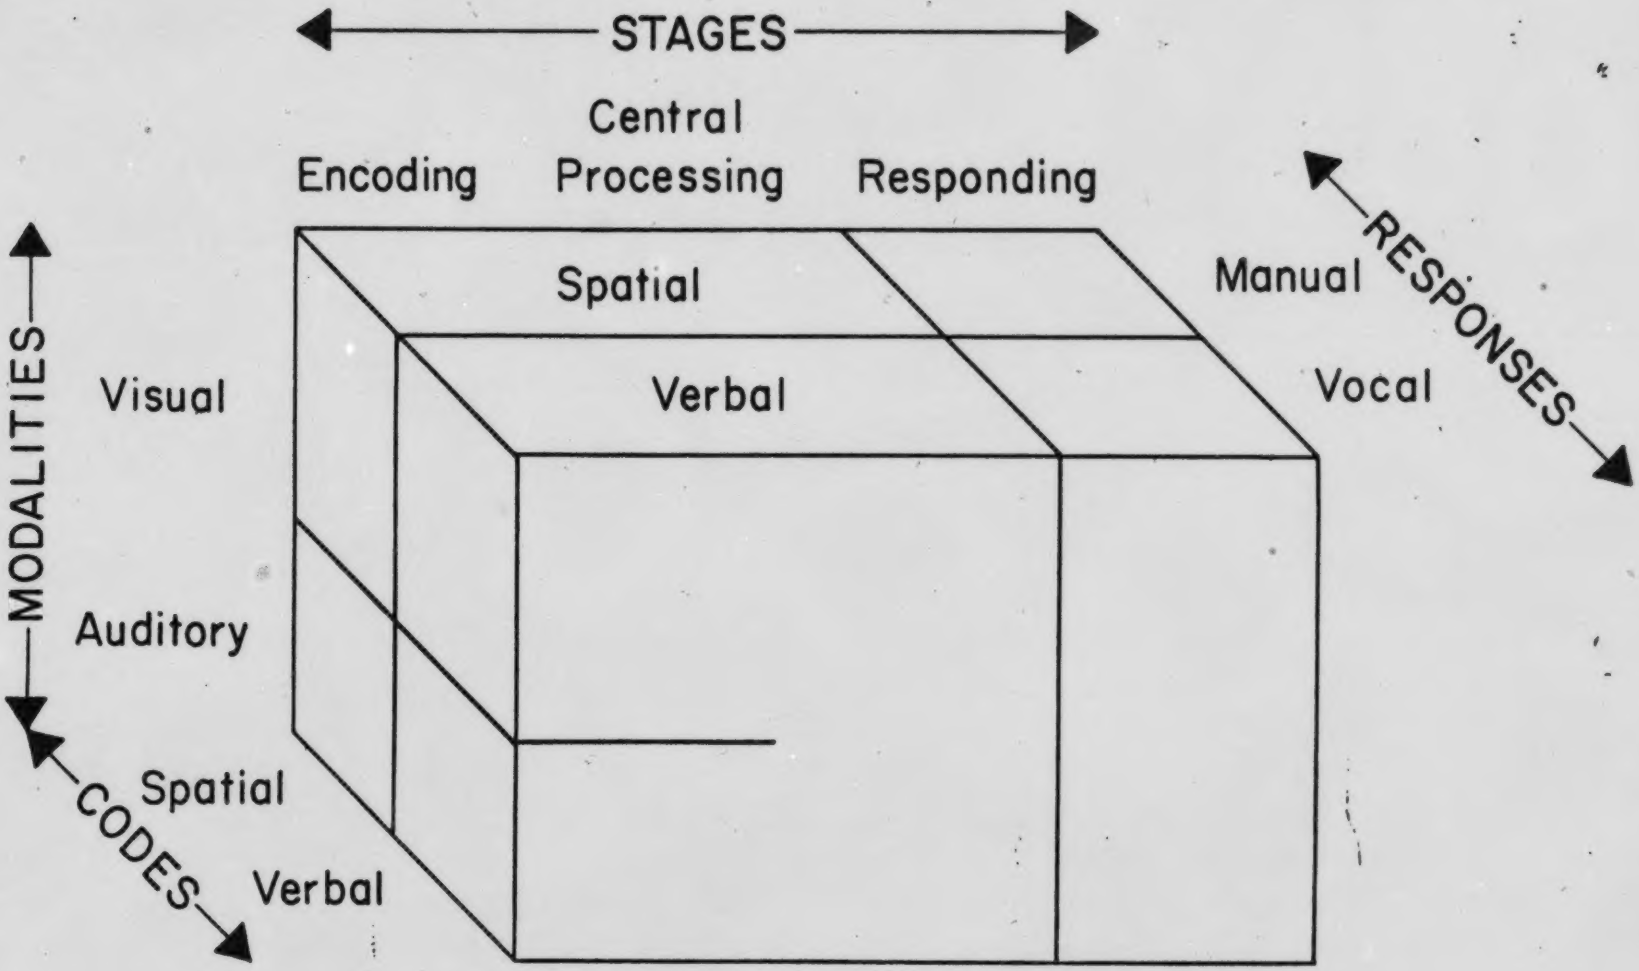
STAGES
Encoding
Central Processing
Responding
MODALITIES
Visual
Auditory
CODES
Spatial
Verbal
RESPONSES
Manual
Vocal
Spatial
Verbal

Thus it is necessary to assume that some capacity is shared across the "pools" of Figure 1. This may be in the form of a single "general" capacity, equally available to all processes, or a (2) more *hierarchical structure* represented in Figure 2. Adopting a hydraulic metaphor, we may assume that resources above a boundary may transfer across the boundary, but those below may not and exclusively serve the indicated process. If this conception is adopted, it becomes important to specify the dominance ordering of the hierarchy. For example, Figure 2 suggests that each *code* (spatial vs. visual) has its own exclusive pool of auditory and visual resources. The implication is that an auditory spatial and an auditory verbal task will be time-shared as efficiently as a verbal and spatial task of opposite modalities.

This prediction runs somewhat counter to our intuitions, which suggest instead (3) a kind of a *"shared features" conception*. Two tasks will interfere to the extent that features (dimensions of resource) are shared between them. Maximum interference will occur when there is complete feature overlap; minimum when there is none, and intermediate levels when one or two features are shared. Such a conception may be represented "hydraulically" in Figure 1, by assuming semi-permeable boundaries separating pools. In this conception one needs again to establish the dominance ordering of the different dimensions (the degree of permeability of the three orthogonal boundaries of Figure 1). In time-sharing perceptual encoding tasks, is it more important that modalities be different between tasks? Or processing codes? Or is there an interaction between the shared features such that, for example, shared modalities make a difference within a code, but not between?

There exists a considerable quantity of experimental data that supports the separate resource-like properties of the three dimensions (Wickens, 1980). However, some of this data is flawed by confounds with task difficulty, (e.g., an intra-modal time sharing combination may involve more difficult task components than a cross-modal), or peripheral interference (e.g., a visual-visual task pairing may induce a degree of visual scanning, not encountered in a cross-modal combination, thereby biasing time-sharing efficiency in favor of the latter). Finally, studies have not extensively examined resource dimensions orthogonally in pairs, or triples.

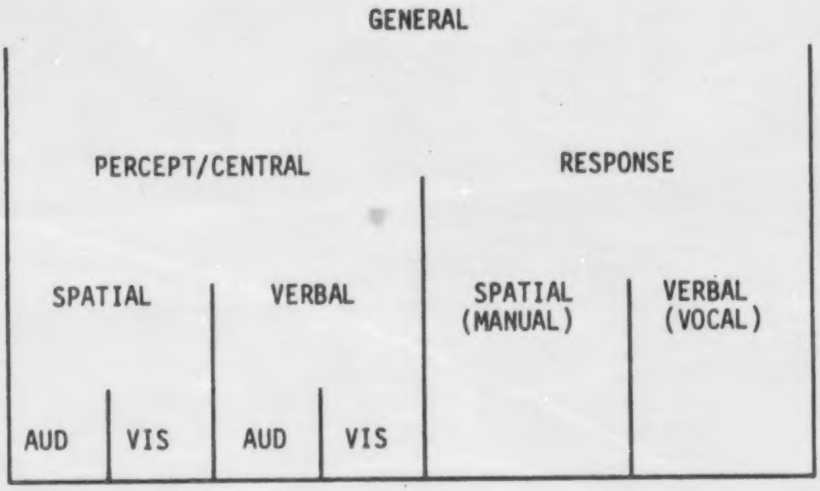
GENERAL
PERCEPT/CENTRAL
RESPONSE
SPATIAL
VERBAL
SPATIAL (MANUAL)
VERBAL (VOCAL)
AUD
VIS
AUD
VIS

Practical Implications

The multiple resources concept has four important implications to human performance in complex settings: These relate to workload, task integration, task structuring, and strategies.

Workload. If resources are a vector quantity and mental workload is conceptualized as the demands imposed by tasks for the operator's resources, then effective workload measures must also be vectors. Alternatively, for workload measures that are scalars (e.g., subjective ratings, physiological measures), their derivation from (or mapping into) the vector resource concept must be specified. For example, if a systems designer derives a subjective difficulty index, is this index tapping perceptual load? Response load? Or come contribution of each?

Task Integration. Most directly the structure-specific resources concept suggests that in high information flow environments (the overloaded ground controller, or jet pilot), greater time-sharing efficiency will be achieved when demands are shifted from over- to under-utilized resource pools. This option is increasingly available, as computer technology has provided great flexibility in display and control options (e.g., voice technology for both display and response capabilities).

Task Structuring. There is good experimental evidence that certain codes or modes of processing are more naturally associated with each other, than others (Greenwald, 1971). For example the manual response system represents a more compatible output from spatial processing than verbal. Thus when a designer considers the task to which for example voice response, or auditory display will be assigned, part of that decision should be dictated by whether certain natural compatibility relations are either confirmed or violated with a particular selection.

Strategies and Training. If resources meet the criteria defined above, then they are capable of differential allocation according to the operator's conscious choice. One implication of the multiple resources model--dealt with in some detail by Navon and Gopher--is that the operator has considerable flexibility in how resources from various pools, are allocated to various tasks. For a given task pair, different allocation policies may lead to varying degrees of time-sharing efficiency, and only one particular choice may optimize performance. Correspondingly, a component task may be performed in various ways that will differentially load different pools (and thereby differentially impact performance of a concurrent task). A short-term memory task, for example, may be performed by emphasizing phonetic versus semantic codes. Signal detection may be varied by changing sensitivity or bias parameters. The supervision and control of dynamic systems may be exercised by emphasizing spatial or propositional (verbal) logic. Recognition of the role of strategies in influencing dual task performance thus has bearing both upon training (the extent to which specific time-sharing strategies may be trained) and upon selection (potential differences in the availability of these strategies between people).

Experimental Results

A large portion of our research on multiple resources has centered around the manual tracking paradigm. This task is advantageous for two important reasons: (1) It is a relatively ubiquitous paradigm that represents a major component of an aviator's task. Yet with only minor changes in certain experimental parameters of the task, the paradigm can be modified to simulate the slower dynamics of ship control, or the more complex transfer functions characteristic of chemical or process control. (2) It is a task that can provide a relatively continuous and rich supply of experimental data, which is able to indicate shifts in the subject's processing strategy. Furthermore, the task is complex enough that a number of different parameters, or chracteristics, can be systematically manipulated, to impose differing loads at all levels of the processing system. Our general experimental approach has been to use well-trained subjects, and allow them to track under single task conditions, and under a variety of dual task conditions during which the difficulty (resource demands) of tracking, and of the concurrent task, is varied. Variance in time-sharing efficiency induced by these manipulations is then employed to make inferences both about the nature of the tracking task itself and about the functional composition of processing resources. This "bootstrapping" operation is substantiated by converging data from other sources.

The following description of experimental results will be organized around four experimental manipulations (some of which were performed in more than one different experiment). These concern tracking order, tracking bandwidth, display modality and hand assignment.

Tracking Order. The order of a tracking task refers to the number of time integrations between a control input (exerted by the subject in response to a perceived error) and the response of the plant. Generally, first order (1 integral) or velocity control systems are fairly easy to control. Second order (2 integral) or acceleration systems are sluggish, unstable, and difficult. Second order dynamics are characteristic of many aviation systems with high inertia. The second order task may be likened to balancing a ball bearing in the center of a flat tray, as an unseen force continuously tilts the tray off the horizontal. Generally, there are two different strategies of coping with second order dynamics. A "perceptual" strategy entails heavy anticipation, as the operator responds directly to higher derivatives (velocity and acceleration) of the error signal. A "response" strategy entails the execution of a coordinated "bang-bang" double impulse control, to nullify an existing error.

Using dual task data from two converging sources, we have found that well-trained operators tend to demonstrate a perceptual-loading strategy. In one investigation, six subjects performed first and second order tracking while evoked potentials to a Bernoulli series of auditory tones were recorded (Figure 3). Subjects covertly enumerated occurrences of the rarer of the two tone pitches. Following procedures described elsewhere (Isreal et al. 1980 a, b), the amplitude of the P300 component of the ERP was examined as an index of the perceptual/central processing load imposed by second order tracking. In fact, the amplitude of P300 declined from no tracking, to first, to second order control, allowing us to infer that the increase in order was one that demanded resources from a pool utilized to generate the P300 (Figure 4). Since P300 was elicited without any response requirement, it is doubtful that this pool was related to response processes. Reaction time, collected concurrently with tracking, in a different paradigm rose in a corresponding fashion with tracking order, validating that the order increase did demand more resources (Figure 4, top).

An alternative interpretation to these data is that resources are undifferentiated, that P300 would be sensitive to any demand manipulation, and

7

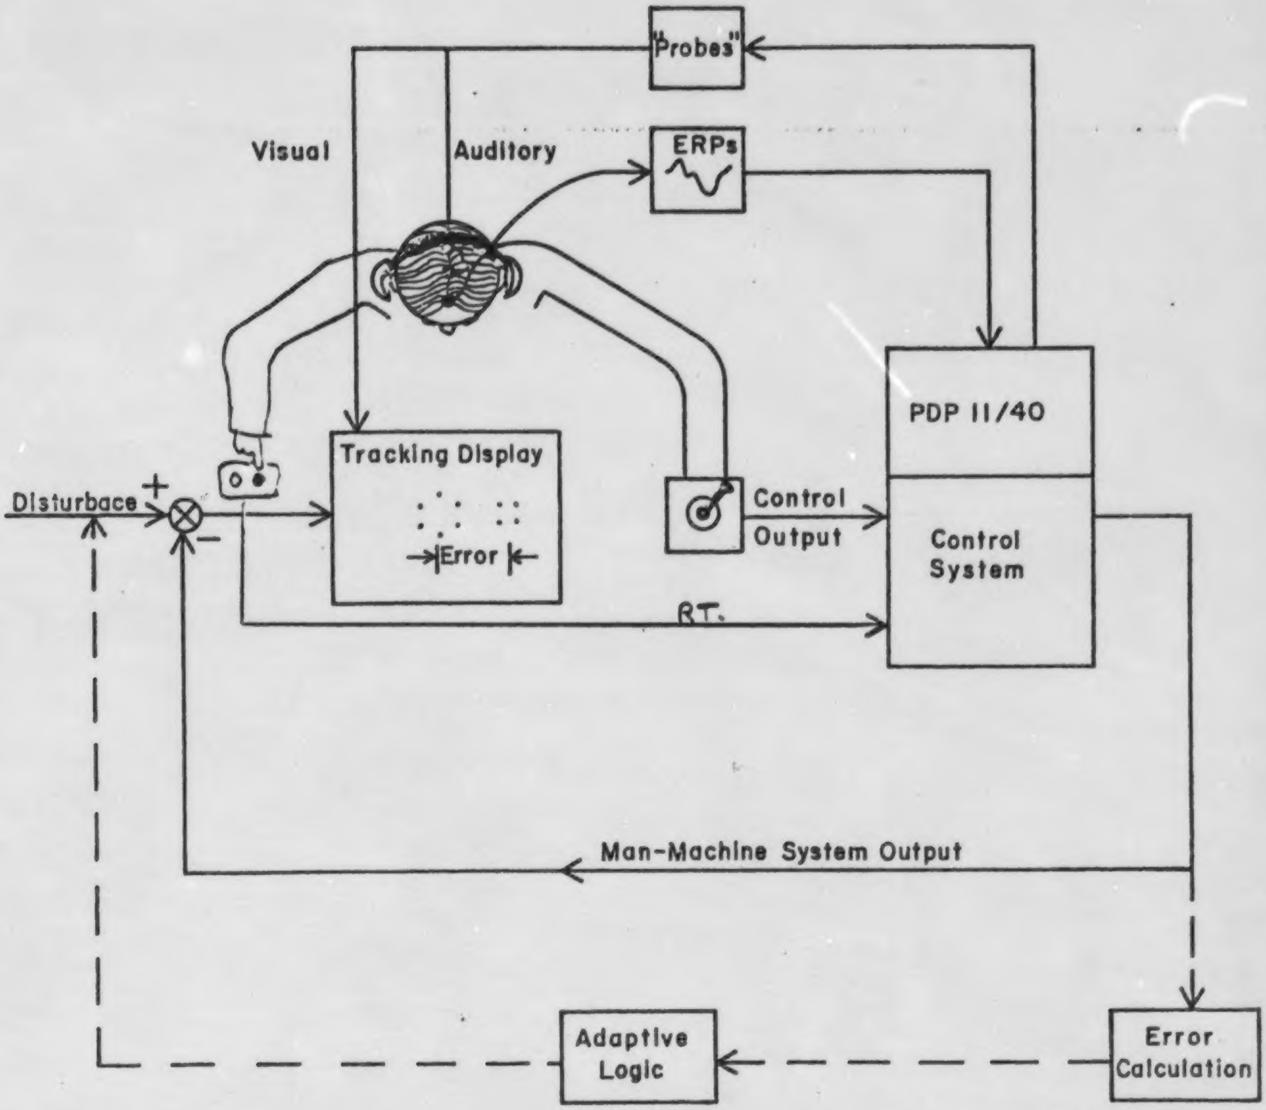
"Probes"
Visual
Auditory
ERPs
PDP 11/40
Tracking Display
Error
Disturbace
+
−
Control
Output
Control
System
RT.
Man-Machine System Output
Adaptive
Logic
Error
Calculation

REACTION TIME

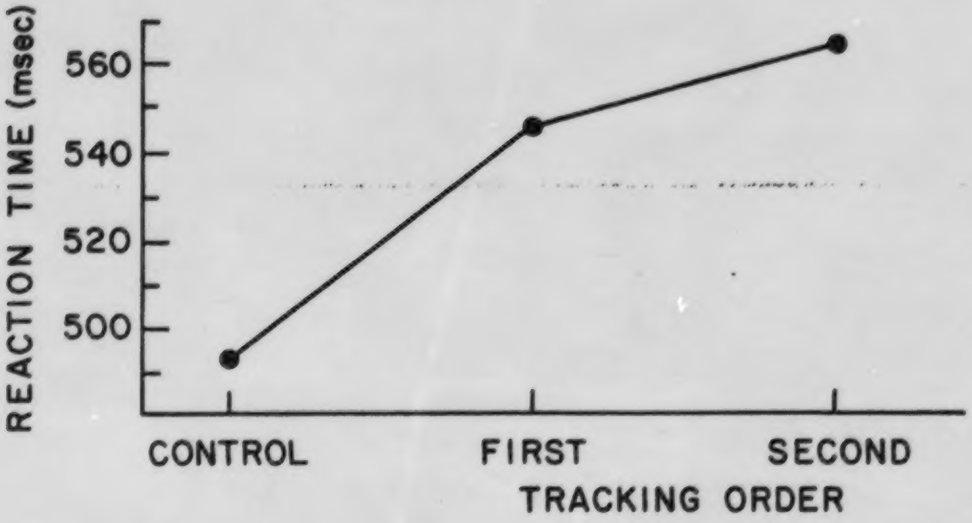

ERP (RECORDED AT P_z)

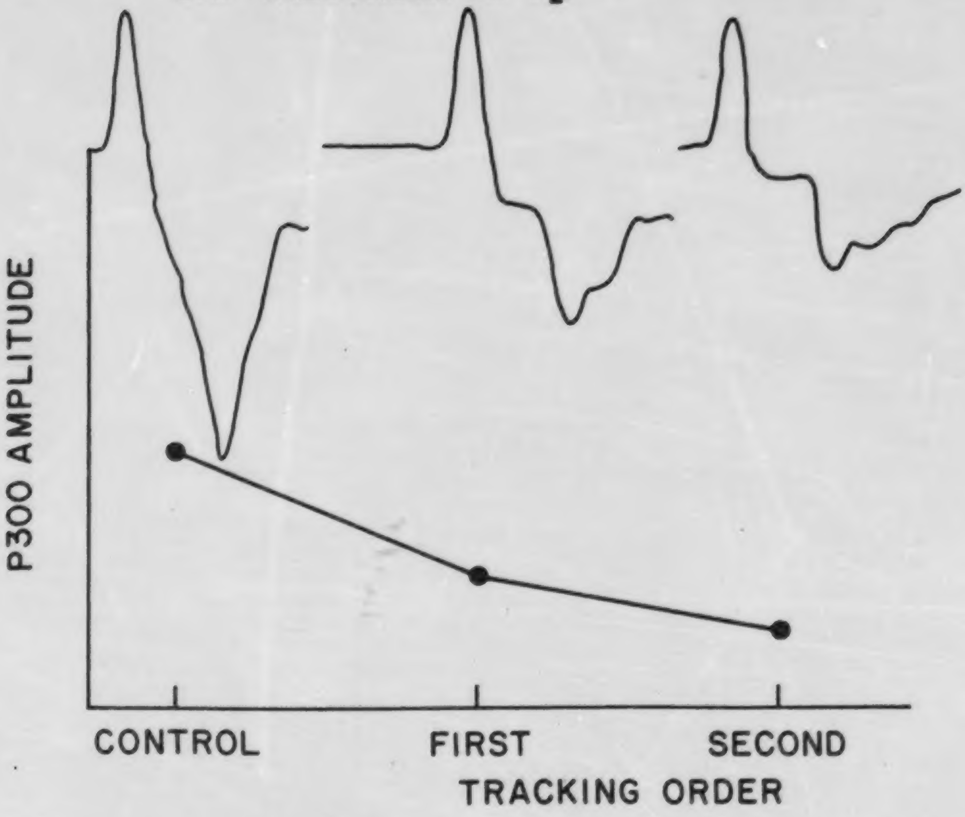

therefore that second order tracking does not necessarily impose a greater perceptual/central load but its load is unspecified. Two lines of evidence, however, argue against this view. One employs a different second task and the other a different tracking task manipulation. These shall be considered in turn.

A second, control-order manipulation experiment was conducted in which a Sternberg memory-search task was imposed as a concurrent task (Wickens et al, 1980). In orthogonal manipulations, control order of tracking was adjusted, and the Sternberg task was rendered more difficult (prolonged) by (1) imposing a display mask, (2) increasing the memory set size (from two to four times), and (3) increasing the complexity of the response (from a single to double key press). The first two of the Sternberg manipulations were intended to demand resources from the perceptual/central pool, and therefore to be more disrupted by any added perceptual load inherent in second order tracking. The double-response load manipulation should be sensitive to any increase in response demands by second order tracking.

The results, shown in Figure 5 (top) confirm that tracking order and manipulations of perceptual and central load interact (reliably in the former case), while order is additive with response load. (In fact, the total absence of effect in the latter case -- at first surprising -- is explainable when tracking error is examined. This variable increases with the double response, but does so to an equal extent with first and second order dynamics.) These results then are consistent with the results of the ERP experiment in locating the demands of second order tracking, as performed by these subjects -- at earlier processing stages.

Tracking Bandwidth

In a series of three experiments, we have manipulated the bandwidth, or upper cut off frequency of the random input forcing function which the subject must nullify in performing the task. Essentially, this manipulation varies the "speed stress" or the number of decisions and motor commands per unit time that must be initiated. Like control order, this manipulation renders the task subjectively more difficult, increases tracking error, and will sometimes generate greater interference with a concurrent task. Yet its effects appear to be qualitatively different from control order.

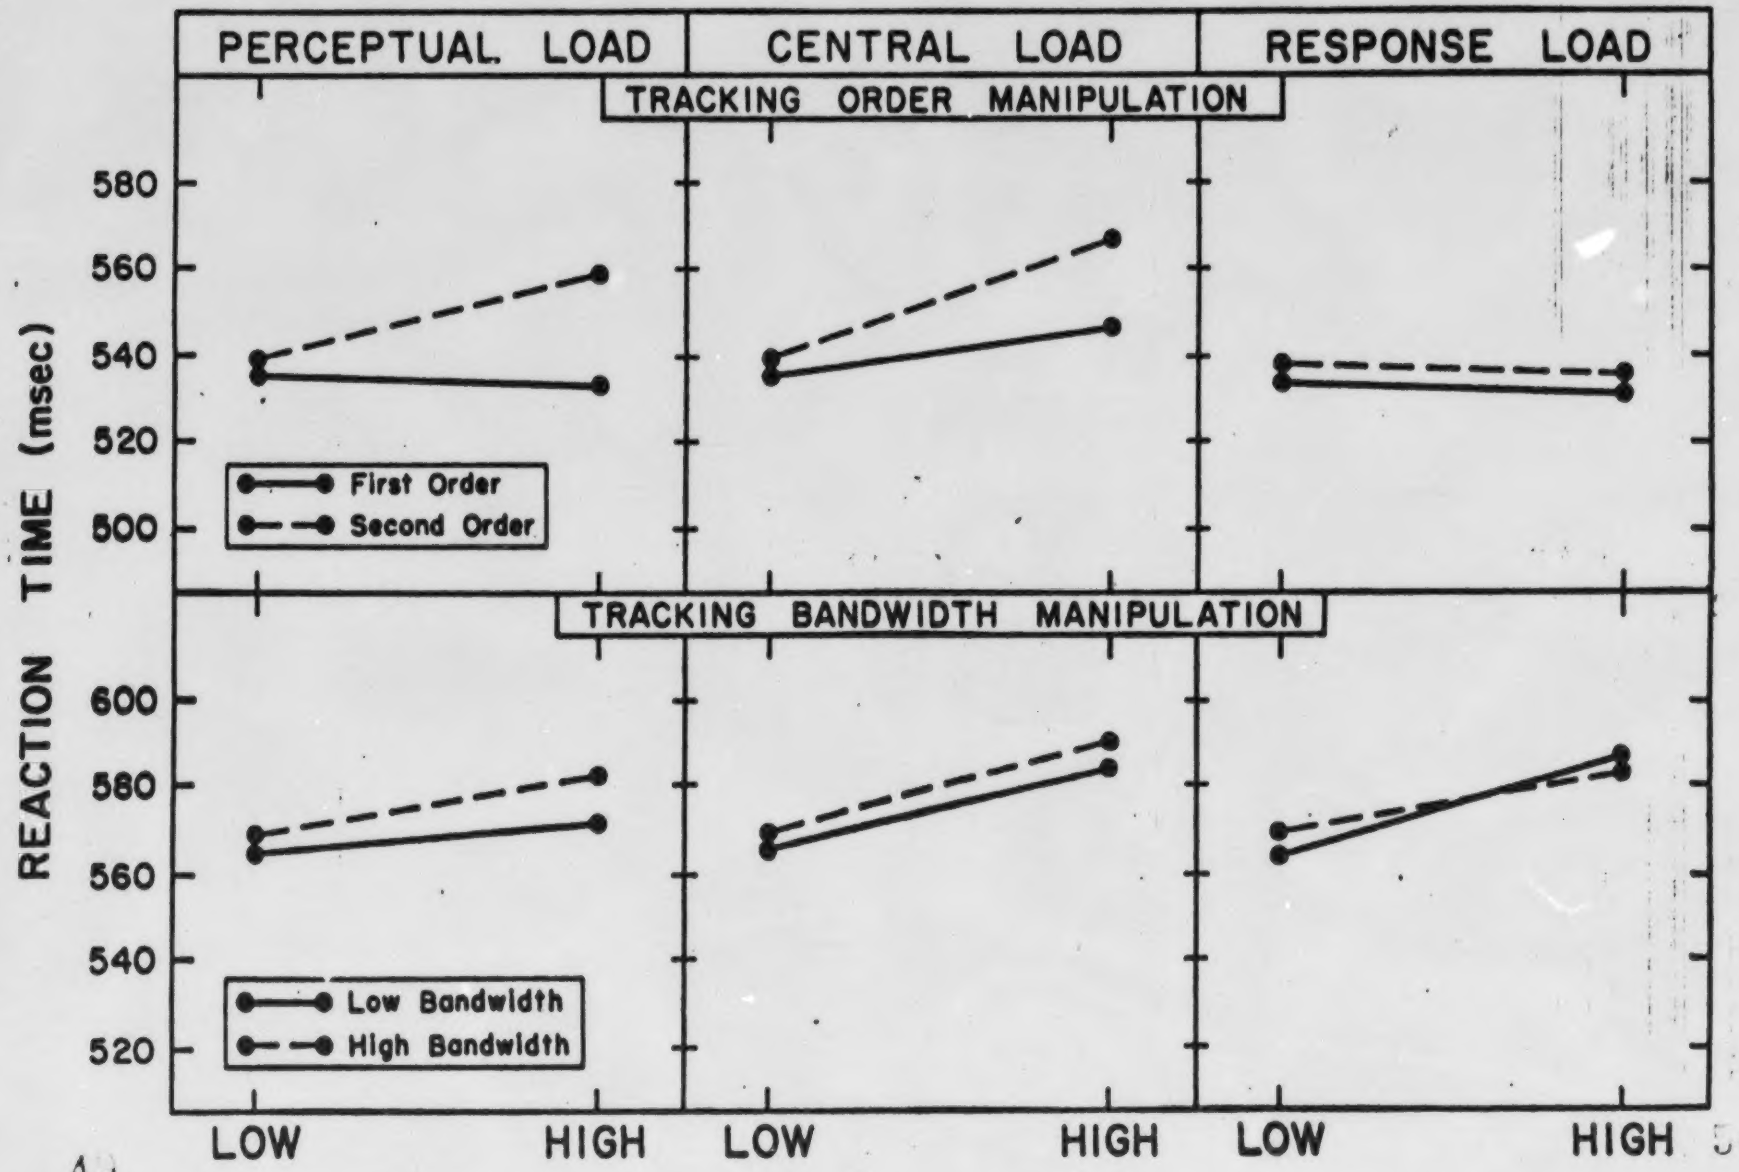
PERCEPTUAL LOAD
CENTRAL LOAD
RESPONSE LOAD
TRACKING ORDER MANIPULATION
REACTION TIME (msec)
580
560
540
520
500
First Order
Second Order
TRACKING BANDWIDTH MANIPULATION
600
580
560
540
520
Low Bandwidth
High Bandwidth
LOW
HIGH
LOW
HIGH
LOW
HIGH

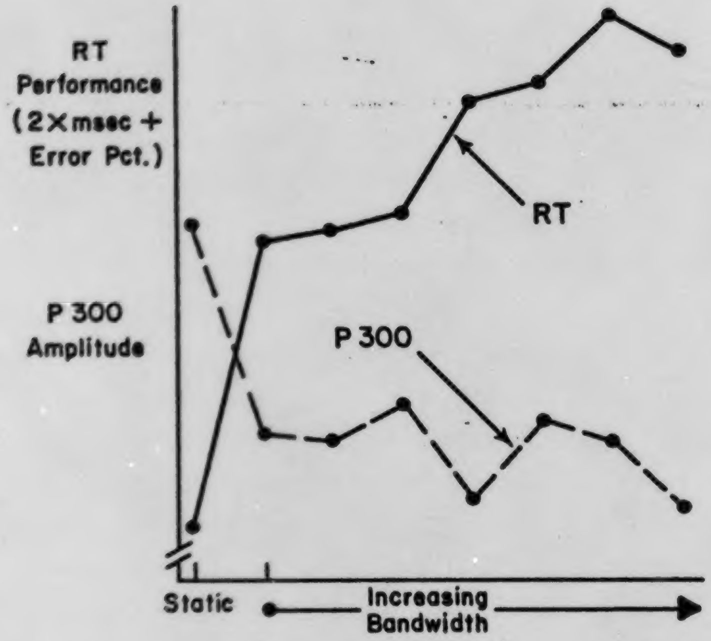
RT
Performance
(2× msec +
Error Pct.)
P 300
Amplitude
RT
P 300
Static
Increasing
Bandwidth

When bandwidth was manipulated[2] concurrently with the ERP and reaction time tasks as above (Isreal et al. 1980a), the results shown in Figure 6 indicate that RT performance declined systematically (and reliably), whereas P300 amplitude was unaffected. These data emphasize that P300 is selective in its resource sensitivity and is not affected by any difficulty manipulation of a concurrent task.

In considering these results, it was hypothesized that the locus of demands imposed by the bandwidth increase was on the response pool of resources. However, the pattern of results of two further bandwidth manipulations suggest that this interpretation may not be entirely correct.

When bandwidth was manipulated concurrently with the Sternberg task (Wickens et al., 1980, Figure 5, bottom), there is no suggestions of an interaction with any processing stage. (The response load hypothesis would suggest that an interaction would be observed with the double response condition.) Furthermore, there is no suggestion that increased bandwidth had any effect at all upon reaction time, despite instructions to the subject to maintain tracking as the primary task, and therefore to cast all variance due to dual task interference into the RT performance.

A similar pattern of results was observed by Wickens and Harris (Wickens, 1980) in an experiment in which bandwidth was manipulated concurrently with a running-memory mental arithmetic task in which subjects computed the absolute difference between successively presented digits. Stimuli for the latter task were presented in either the visual or auditory modality, and responses were generated either manually or vocally. The results (Figure 7) expressed in terms of a total dual task decrement measure indicated (a) no main effect of bandwidth (that is, increasing bandwidth did not increase the competition for resources with the arithmetic task) and (b) no interaction of bandwidth with input or response modality. Had the locus of bandwidth increase been at response, we might have anticipated a greater effect in the manual, as opposed to the vocal, response condition, since in the former condition the output channel is shared by the the two tasks.

---

[2] In fact, a continuous range of bandwidths was assayed, bandwidth being modulated upwards and downwards over time.

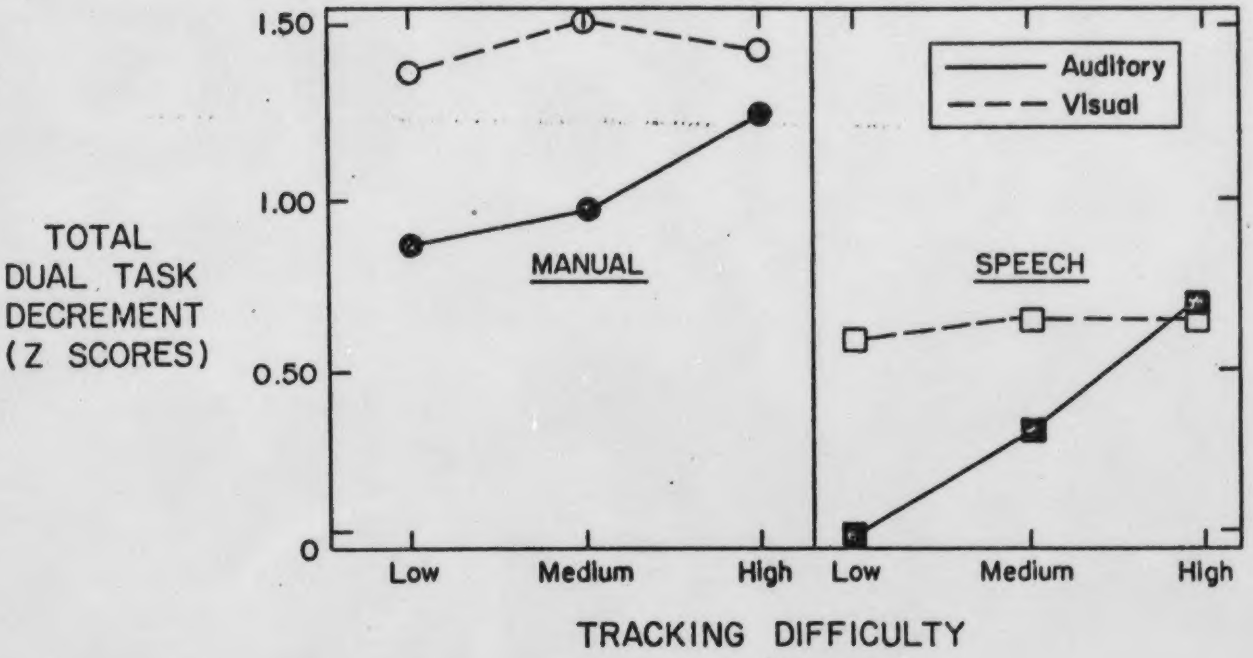
TOTAL DUAL TASK DECREMENT (Z SCORES)
1.50
1.00
0.50
0
MANUAL
SPEECH
Auditory
Visual
Low
Medium
High
Low
Medium
High
TRACKING DIFFICULTY

The results of the latter two studies were somewhat surprising. Collectively they suggest that the resources consumed by increasing bandwidth are not utilized in the Sternberg or mental arithmetic task. One possible clue to the interpretation is provided by the fact that both of these tasks are verbal in their processing requirements (encoding and searching letters of the alphabet, and retaining and subtracting digits respectively). The RT task employed in the ERP bandwidth experiment was non-verbal, requiring only the pitch discrimination of two tones. This might suggest that the resources underlying bandwidth are spatial in nature (an increase in the rate at which spatial uncertainty needs to be resolved). This consideration leads to a discussion of two further experiments: one related to the relation between the spatiality of tracking and processing modalities, and the second to the relation between spatiality and hand of control.

Display Modality

If tracking truly requires spatial processing, then it might be expected to interfere to a greater extent with a second task that also requires spatial processing, than with one that does not. A related issue concerns the relation of spatiality to modality. Normally the visual modality is associated more directly with spatial processing, and most demonstrations of spatial interface occur with visual tasks (but see Baddeley and Lieberman, 1980). An experiment by Isreal (1980) therefore asked whether the spatial aspects of tracking could be unconfounded from the modality aspects, to establish if a common spatial processor transcends across modalities.

In Isreal's experiments, nine subjects performed a one-dimensional compensatory tracking task that was displayed either visually or auditorily. In the latter condition, "error" was displayed by a tone that varied redundantly in pitch (proportional to absolute error) and apparent spatial location (proportional to signed error). The subject therefore tracked in such a way as to maintain the tone in the mid plane of the head, and at the lowest possible pitch. Because of the lesser degree of familiarity with this display, subjects were provided twice as much practice as with the visual task, and the auditory bandwidth was adjusted to a lower level so that RMS error obtained on the two displays was equal.

Concurrently, subjects performed a disjunctive go/no-go reaction time task, presented in either the auditory or visual modality. This required either a spatial discrimination (of a tone in the left or right ear, or a bar on the left or right side of the display), or an intensity discrimination (a low intensity or high intensity bar flash, or a soft or loud tone). Stimulus discriminability was adjusted so that single task RT in all conditions was equivalent but for a 40 msec difference favoring the auditory modality. This difference accounted for the greater peripheral transmission latency of the visual modality.

Figure 8 presents a measure of total dual task interference (combining decrements in RT latency and accuracy and tracking error from their respective single task values). A number of reliable trends are evident from the data. (1) There is an overall "spatial effect" since the decrement with the spatial probes is enhanced relative to that of the intensity discrimination probes. (2) This spatial effect is enhanced when probes appear within the same modality as the tracking task (the four outermost probes of the figure). (3) There remains a reliable cross-modal spatial effect (the four inner points). (4) Intramodal interference is greater in the auditory than the visual modality, despite the equivalence of single task tracking performance. The detailed interpretations of these results are presented elsewhere. However, they have some important implications both with regard to tracking and to the structure of resources. The confirmation of the spatiality of tracking is perhaps not surprising, but it is important to note that this effect is unequivocal even when all other characteristics of the spatial and intensity probes were equated (including their response component -- a non-spatial press of a single button). The cross modal spatial effect confirms the existence of modality-free spatial processing (Baddeley and Lieberman, 1980), but its enhancement within the modalities suggests a modality-specific component to spatiality as well. Concerning a dominance relationship between shared modalities and shared spatial processing as outlined earlier, the answer appears to depend upon the tracking modality. When visual tracking is employed, the spatial effect clearly outweighs the modality effect. Yet with auditory tracking, the relation appears to be reversed.

These latter results suggest an asymmetry either between tracking tasks, or modalities. It is possible that the auditory task -- despite its initially

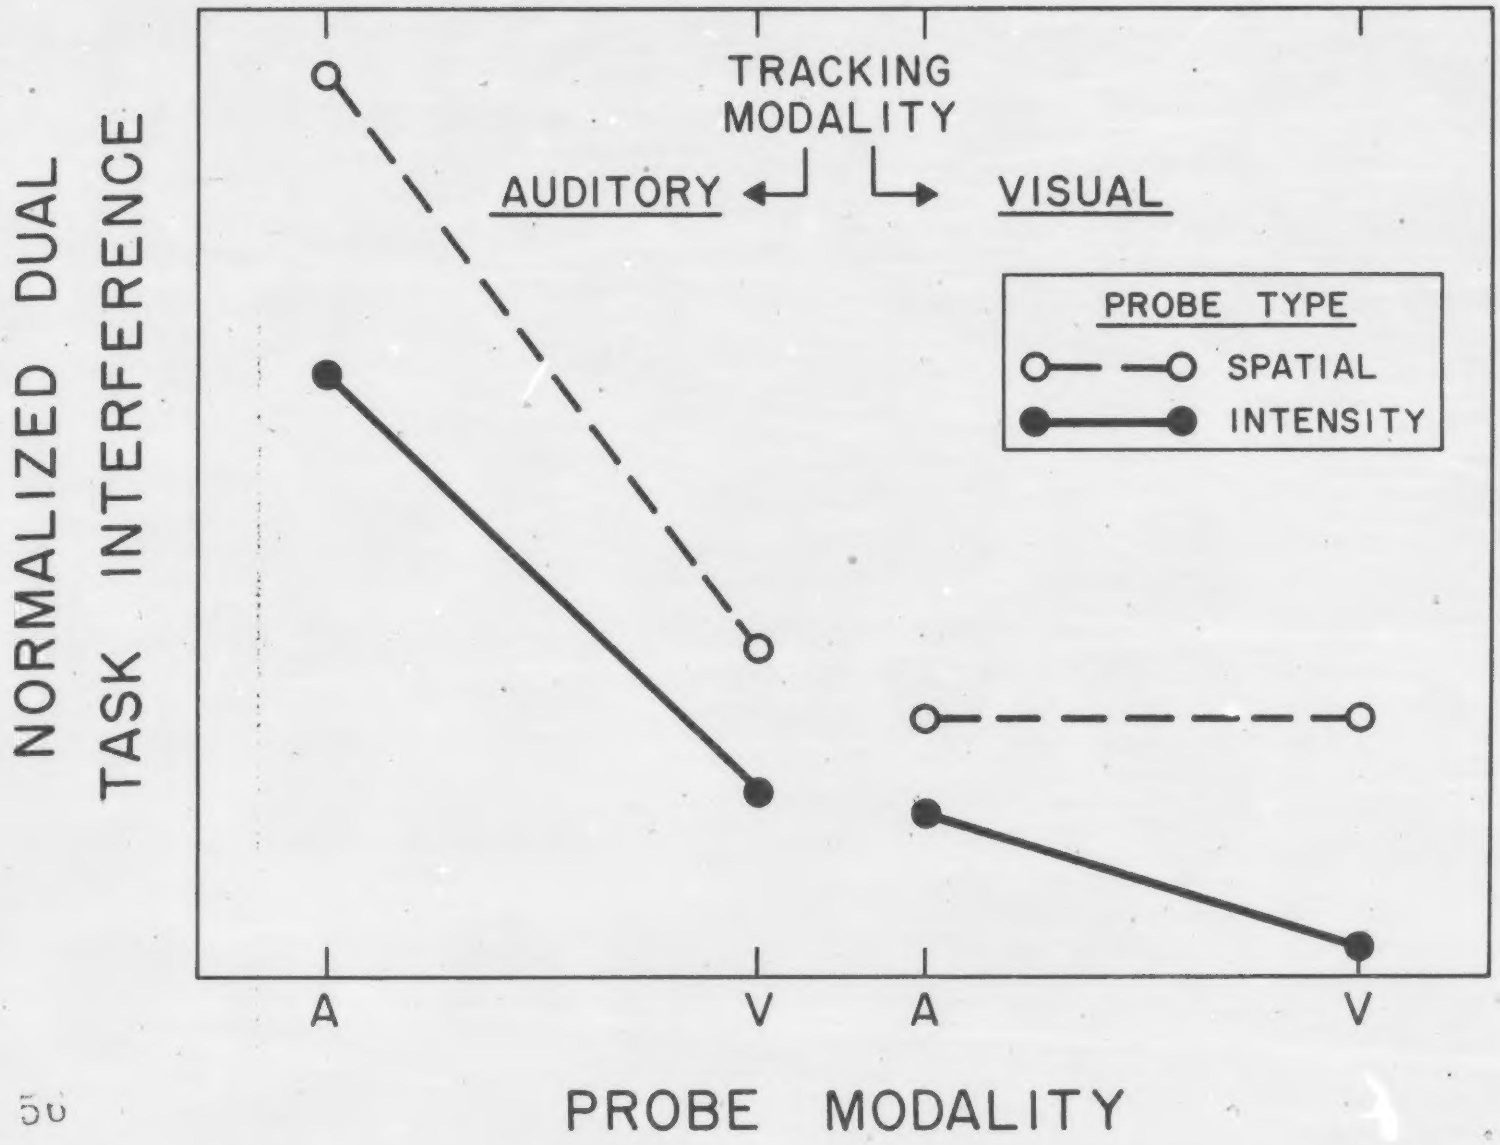
TRACKING MODALITY
AUDITORY
VISUAL
PROBE TYPE
SPATIAL
INTENSITY
NORMALIZED DUAL TASK INTERFERENCE
A
V
A
V
PROBE MODALITY

greater practice, easier level and equivalent performance -- has a greater resource limited region than the visual task and is therefore susceptible to greater interfering effects. This is due perhaps to a less natural association between that modality and spatial processing. Alternatively, the auditory modality itself may have less time-sharing capacity, independent of the task characteristics. This interpretation is consistent with a similar trend to that observed in Figure 7, when Treisman and Davies' (1973) data from an analogous experiment with verbal material are replotted in an equivalent manner. A related finding, also consistent with Treisman and Davies' results and those of Hunt and Lansman is that in cross-modal conditions, the auditory task always suffered the greater decrement, independent of whether this was the tracking or the reaction time task.

Hand Assignment

The previous study went further to suggest the spatiality of tracking, but did nothing to associate that spatiality with the possible residence of spatial processors within the right cerebral hemisphere, as suggested by a wealth of experimental and clinical data. Indeed, functionally it may make little difference where spatial processing "resides" as long as its resource demands are somewhat disjoint from verbal processing. Nevertheless, the data from the fourth experiment (Wickens and Sandry, 1980) suggest some hemispheric lateralization of the tracking task, and go further to indicate a potential advantage in time sharing realized when certain codes of processing, or resources, are associated with each other in a compatible relation.

In this experiment, we proposed a condition of "task-hemispheric integrity" to exist when the processing and response functions of a single task are carried out within a single cerebral hemisphere (e.g., a spatial processing task is responded to with the left hand, and a verbal one with the right hand, or with a vocal output). The hand assignment that generates integrity may be unlikely to facilitate single task performance (and indeed may hinder it), but under dual task conditions when a verbal and spatial task are shared, this assignment is predicted to maximize efficiency, because the processors within each hemisphere are only involved with a single activity.

Eight subjects tracked with either the left or right hand alone, and concurrently with a Sternberg Memory Search Task. The latter was also responded with each hand in turn. Single and dual task performance on each was assessed as a function of hand assignment. The basic comparison contrasting dual task efficiency in the integrity condition (tracking left, Sternberg right), with the non-integrity condition, revealed the former, with practice, to be reliably superior. A second experiment replicated this design with a "spatial" alphabet of 26 random dot patterns. Since both tasks here were presumably spatial, little difference due to hand assignment was expected, and none was in fact observed. Tentatively, the results of this study provide the spatiality presumed to underlie processing in the tracking task with an anatomical residence, by virtue of the assumed ipsilateral control of respond hand.

Conclusion

The experimental data reported here only begin to address the complexity of tracking demands and of the resources underlying human information processing. In particular, these have not addressed the potential role of strategies alluded to above either as these alter performance of a task (e.g., employing a response strategy with second order tracking to reduce interference with a cognitive task, or using verbal-propositional representation of complex system dynamics to avoid interference with a spatial task), or strategies are used to adjust the resource allocation policy. The research of Gopher and his colleagues suggests that important benefits to dual task performance can accrue with training on this adjustment. Also not addressed has been our work on the implications for task workload assessment.

I believe that the data presented here, along with the converging evidence from numerous other laboratories, support the assertion that a multi-dimensional model of resources can account for a large proportion of variance in laboratory data from time-sharing experiments, as well as an understanding of the way in which manual control tasks are performed. The important issue that remains, however, is the extent to which these data have a bearing upon the pilot's performance in the air. Does he truly time-share? Or, in fact, is processing typically "single channel" in the real world (Moray, 1980). Alternatively, does the proportion of variance accounted for by structural differences reduce to a fairly trivial level, in relation to variables such as time pressure, or task complexity (e.g., the demands within a pool)? As so often is stated, answers to these questions await empirical validations.

References

Baddeley, A. & Leiberman, K. Spatial working memory. In R. Nickerson (Ed.). *Attention & Performance VIII*. Erlbaum, 1980.

Greenwald, A. Sensory feedback mechanisms in performance control. *Psychological Review*, 1970, *77*, 73-79.

Isreal, J. Structural interference in dual task performance: Event-related potential, behavior, and subjective effects. University of Illinois, Unpublished Ph.D. Dissertation, 1980.

Isreal, J., Chesney, G., Wickens, C. & Donchin, E. (a) P-300 and tracking difficulty: Evidence for multiple resources in dual task performance. *Psychophysiology*, 1980, *17*, 57-70.

Isreal, J., Wickens, C., Chesney, G., & Donchin, E. (b) The event related brain potential as an index of display monitoring workload. *Human Factors*, 1980, *22*, 211-224.

Kahneman, D. *Attention and Effort*, Prentice Hall, 1973.

Moray, N. Attention in Process Control. In W. Rouse & J. Rasmussen, (Eds.). *Human Detection and Diagnosis of System Failures*. N.Y.: Plenum Press, 1981 (in press).

Navon, D. & Gopher, D. On the economy of the human processing system. *Psychological Review*, 1979, *86*, 214-255.

Norman, D. & Bobrow, D. On data-limited and resource-limited processes. *Cognitive Psychology*, 1975, *7*, 44-64.

Treisman, A. & Davies, A. Divided attention to ear and eye. In S. Kornblum (Ed.). *Attention and Performance IV*. N.Y.: Academic Press, 1973.

Wickens, C. D. The structure of attentional resources. In R. Nickerson (Ed.). *Attention and Performance VIII*. Erlbaum, 1980.

Wickens, C., Derrick, W., Micalizzi, J., & Beringer, D. The structure of processing resources: Implications for task configuration. Proceedings, 1980 Annual Meeting of the Human Factors Society, 1980.

Wickens, C.D. & Sandry, D.L. An investigation of the dual task performance relationship to hemispheric specialization. University of Illionois, EPL Technical Report EPL-80-1/ONR-80-1. June, 1980.

A COMPUTATIONAL THEORY OF THOUGHT APPLIED TO TRADITIONAL PROBLEMS

EARL HUNT

THE UNIVERSITY OF WASHINGTON

Thinking is the act of manipulating an internal representation of the external world (Johnson-Laird, 1980; Newell, 1980). The goal of cognitive science is to discover the form of this representation. Thinking is also an act that must be done by something. Thus a cognitive science theory of thinking must meet two requirements; it must be an accurate reflection of the environment as seen by the thinker, and it must be a symbol system that can be realized by the thinker's brain, a particular physical organ. Both requirements restrict our thinking, but opinions differ markedly concerning the relative importance of each type of restriction. At one extreme, the computer simulators argue that if we can find just one symbol system that can solve the problems that people do, then that system is immediately a candidate psychological theory (Newell, 1973). This implies that the environment is extremely restrictive, and that what we should do is examine invariances in our response to it. An alternative view is that the environment can be represented in many ways. If this is true, it will be more efficient to determine what basic functions the brain uses in constructing a representation. Most of experimental psychology follows this tradition. Finally, physiologically oriented psychologists argue that the physical nature of the brain restricts the possible representations it can construct. If this is true, cognition can be understood only after we understand the physiology of the brain.

Our beliefs about the importance of each of these restrictions determines what sort of thinking we choose to study. Cognitive science studies, growing out of artificial intelligence and linguistics traditions focus on invariances, while experimental psychology focuses on the relation between changes in the environment and changes in performance. While, in principle, the physiological approach could be applied to the study of invariance or change, in practice change is emphasized. Changes in thought that are systematically associated with changes in the thinker have been of particular interest, for thinking cannot be described fully unless we know whose thinking is being described. Individual differences in cognition are far from random; they are systematically related to non-psychological attributes such as age, sex, and education in the broadest sense. Furthermore, intra-individual changes in thinking can be produced by temporary changes in physical status, e.g. by fatigue, time of day, or drug state.

What has been learned from these different approaches? The Cognitive Sciences have produced a number of case studies (simulations) and some loosely stated general principles that are said to be dictated by these studies. Experimental psychologists have produced detailed models of how people respond to changes in specialized laboratory environments, but have not combined these models into a unified theory of mental action. Studies of individual differences have produced "theories of intelligence" that are more schemes for classifying mental behaviors in terms of their correlates than they are models of the thought process itself. The theoretical contribution of physiological studies of cognition is even less clear. With the possible exception of Luria's (1966) attempts to develop a model from neuropsychological evidence, the literature provides us with data that a theory must incorporate, but no clear picture of how the incorporation is to be achieved.

PROGRAMS, LANGUAGES AND MACHINES AS THEORIES OF COGNITION

Is there a basic assumption about cognition that may unite diverse approaches to mental phenomena? The Cognitive Science approach does contain a coherent assumption about the nature of a cognitive theory. The assumption, which will be called the 'computational model' approach, is that cognitive behavior is not explained unless one can model the behavior by a sequence of computations. The assumption is not equivalent to computer simulation, although the two are related. Pylyshyn (1980) has pointed out that computational models exist at several levels. At the lowest level there is the program intended to simulate behavior in a particular situation. At progressively higher levels we have the language in which the program is written and the machine on which it is executed. In most cognitive science studies considerable attention is paid to the language, as it is supposed to define a virtual machine that embodies general principles about cognition. The physical machine is treated as a convenient computational aid, without theoretical meaning. This is certainly true when we deal with conventional computers. Machine considerations cannot be so lightly dismissed if we wish to deal with brain-behavior relations.

Figure 1 represents the relationship between programs, languages, and machines and, more importantly, between each of these levels and the different types of influences upon cognition. Obviously, a program is written in a language, and a language is executed by a machine. We represent information in the environment by input to a program, and we represent cognitive behavior by the program's output. Only the program is tied to observables at both the input and the output stage. Physical variables can be represnted by machine specifications, but their influence will be filtered through the program, just as output of the program will be modulated by machine specifications. The programming language is not tied directly either to input or output, a point which has substantial implications for the logical status of programming languages as psychological theories.

....

FIGURE 1 HERE

....

If the cognitive sciences effort is successful we will eventually produce a library of programs. These will be written in different languages, and there will be some machines (e.g. sufficiently large computers) capable of executing any of the programs in any of the languages. This situation is depicted in Figure 2. There are two 'dreams' of the cognitive science movement that must be dispelled at this point. One is that every running program written in some psychological language will be an acceptable model of human behavior. This seems extremely unlikely. The programs that are starred in Figure 2 represent programs that are stated in a simulation language but are not, themselves, realistic simulations. The other dream is that there will be just one programming language that can generate realistic simulations. If this were true, the constructs of this language could be assumed to have psychological reality. The Production notation (Newell, 1973; McDermott and Forgy, 1978), which will be discussed in some detail later, has been proposed as a candidate language.

The assertion that a particular language is the appropriate notation for psychology cannot be tested by pointing to examples of its use. There are languages whose superficial structure, at least, is decidedly different from a production language, but that can be used to write any simulation program that has ever been written. The inherently recursive General Problem Solver program (Newell and Simon, 1963) was rewritten in FORTRAN (Quinlan and Hunt, 1968), and I strongly suspect that any other simulation program could be rewritten in a similarly primitive language, if one wished to take the trouble.

Figure 2 seems to more accurately depict the goal of the present Cognitive Science approach. Good and bad simulations will be written in a variety of computing languages.

....

FIGURE 2 HERE

....

How did the chaos of Figure 2 arise? The problem is that computing languages sufficient to write simulations (and, in particular, the production notation) are powerful enough to be equivalent to universal Turing machines. Other, non-psychological, languages have the same power, so there is no way to contrast languages as psychological theories on the basis of the behavior of programs written in those languages. Faced with this situation, the usual appeal is to the "naturalness" of a program written using the constructs of one language or another. Production matching is said to be a natural psychological process, whereas Fortran subroutine calls are not. This is clearly not an objective evaluation.

Pylyshyn (1980) proposed an interesting solution to this problem. He would begin with an excessively powerful language and progressively restrict it until it generated only realistic simulation programs. Furthermore, the restrictions should be justifiable in terms of our knowledge of psychological limits on human thought. Such restrictions are said to be 'principled'. The idea may be illustrated by considering the production system notation itself. A production is a rule of the form

(1) P -> A;

where P is a pattern that may be recognized, and A is an action to be taken. The usual interpretation is that productions are resident in long term memory, and that a production is activated when its pattern is exemplified in working memory (Newell and Simon, 1972). A principled restriction would be to require that P contain at most nine symbols, in conformity with the now classic observation that this is an upper limit on short term memory (Miller, 1956).

Pylyshyn's suggestions are difficult to follow so long as one is confined to a linguistic approach. The sorts of principled restrictions that seem to be needed refer more to the characteristics of the machine on which a program is being executed than to the characteristics of the language in which the program is written. The restriction on the size of productions is a good example. It sounds like the restrictions placed on, say, the level of nesting of parentheses, that might arise from a particular implementation of a general language. A second problem arises when we study inter and intra-individual differences. Under what circumstances do we change the program, the language, or the principled restrictions? Sometimes this issue is easy to handle, when we observe different strategies that people use. The theoretical issue is more difficult to resolve when physical alterations in the brain produce changes in behavior that would be determined at the program level in a simulation. Do we want to say that an executive program in the brain of a deep dyslexic patient (Morton and Patterson, 1979) suddenly switches programs for speech comprehension following a cerebral insult?

Figure 3 summarizes this argument, by presenting a picture of a more realistic "target state" for the cognitive science effort. The figure shows a restricted machine that is designed to execute some programs in a powerful language. Any program in the language, that can be executed by the machine, would be a simulation. Other programs written in the same language, for other machines, would not be simulations, although they might be of interest as examples of Artificial Intelligence. The language alone would be too powerful to be a psychological theory. If the machine's design is adequately motivated by principled restrictions, the implementation of the language on the machine would be a psychological theory, a computational theory of thought.

....

FIGURE 3 HERE

....

In the following section a computational theory will be presented, and then used to describe some of the major phenomena studied by experimental, physiological, and psychometric psychologists.

68

A PRODUCTION EXECUTING MACHINE

The computational theory to be presented consists of a language, the well known production notation, and a machine for implementing productions, subject to principled restrictions. Various bits and pieces of the machine have been proposed in previous papers (Hunt, 1966; 1971; 1978; 1981; Hunt and Poltrock, 1974). Similar proposals have been made, in different terms, by Anderson (1976; Anderson and Bower, 1973; Anderson, Kline and Beasley, 1978) and the relation to the work of Newell and Simon (1972) will be clear. For brevity, I shall refer to a computational theory machine, CTM. Although this phrase is somewhat clumsy in English, it does capture the idea that the machine is only part of a more general theory of thought.

The CTM is based on the idea that thinking can be modeled by programs written in the production notation, and that this notation has psychological significance (Newell 1973; Hunt and Poltrock, 1974). Productions are thought of as rules resident in long term memory (LTM). Productions are executed by comparing their pattern parts to the contents of working memory and, if a match is found, taking the appropriate action. Since one of the actions that is always taken is to change the state of working memory, a "thought", in a computational theory, is a succession of states of working memory.

The production notation can be used to express complex thought sequences by tying productions together into "production systems", sets of productions required to solve a particular type of problem, such as a chess puzzle. As a simple example, consider the following system for manipulating algebraic statements. The system contains the rules

R1. -(-X) -> +X

R2. XY + XZ -> X(Y + Z).

In these rules X, Y, and Z refer to any well formed algebraic expression. Now suppose that Working Memory contained the expression

E1. AB-(-AC).

By application of rule R1 -(-AC) would become +AC. The pattern for R2 would be satisfied, producing the expression A(B+C).

FIGURE 4 HERE

Figure 4 depicts the basic steps that must be taken by any production executing machine. The first step is pattern recognition. Each production's pattern part must be matched against the current situation. This is not necessarily an either-or decision, there can be degrees of matching. If more than one production has a pattern that appears to be matched, the machine must decide which production is to be executed. This is called the 'conflict resolution step'

(McDermott and Forgy, 1978). Conflict resolution must be carried out without considering the 'context of the situation', i.e. the implications of choosing one production over another given the current representation of the environment. The reason for this is straightforward. If the machine considers the current context, we have to explain the psychology of a homonculoid executive program. Finally, the machine must be able to execute the commands for action contained in its productions. The action part of a production must be stated in a vocabulary that refers to actions the machine can take.

Figure 5 depicts the architecture of a machine that consists of the following parts:

FIGURE 5 HERE

----

a. A sensory system connecting the machine to its environment

b. A long term memory system that contains the information that the system has about its past. This information is stored in the production notation. Each "elementary" piece of knowledge that the machine has is expressed as a production. The organization of productions into systems will be considered in more detail in a moment.

c. A working memory system that contains information about the current situation. Just what working memory should contain is a matter of considerable debate. The commonest view is that it contains sections that, loosely, correspond to the sensory systems. I would like to take the somewehat vaguer position that working memory contains an interpretation of the external world that depends upon perceptual responses that have been initiated by sensory input, but are themselves represented as productions resident in LTM. This argument is developed more fully in Hunt (1978).

d. An operating system containing a set of mental operators defining the actions that the system can take. These operators can be thought of as defining the vocabulary in which the action (A) parts of productions can be written. The set of mental operators would have to include operators for altering new production in LTM. Anderson et al. (1978) provide some suggestions. Commands to external effector systems would be part of the mental operator set. Neither learning nor motor performance will be discussed in any length here.

e. A decision system that carries out the conflict resolution step.

Figure 5 illustrates the flow of information and control between these systems. Information from the sensory system and from Working Memory address LTM directly. "Broadcast" than "Address" would be more appropriate because the input is not directed to any particular location. Selection of information takes place within LTM, as a production will simply fail to recognize any input that is random with respect to its pattern.

Working memory receives its input directly from the operation system. That is, every input into working memory (and every change of the contents of working memory) is assumed to result from the activation of one of a limited number of basic mental operations. These operations can be looked upon as the computing circuits of the CTM. The action part of a production must be a command to execute one or more basic operations using specified data in LTM and working memory as input. Newell (1980) has provided a list of some of the operations that any symbol manipulating device would require, e.g. transferring information between working and long term memory. Simple transformations of stored records or simple operations involving the merging of data may also be executed by the basic operations. For example, we combine color, location, and shape information in order to see unitary visual percepts. Thus in a computational theory the 'laws of perception' provide one clue about the nature of the basic mental operations.

The operations, in turn, can receive two sorts of inputs from Long Term memory; a control signal indicating that a particular operation is to be activated, and a data signal indicating the input on which it is to act. Thus the input from long term to working memory is not analagous to simply reading a record, although this presumably does occur. Basic transformations of stored records, or simple operations on a few records, may also serve as commands to this system.

Conflict resolution takes place in the control system. This system receives at most one input from each production in LTM, a signal indicating the extent to which that production has been able to match its pattern part to the signal currently being broadcast throughout LTM. The strength of the signal from a production will be a joint function of the correlation between broadcast signal and the pattern part of a production and of the current 'activation threshold' of that production. A formalism is useful to express this point. Let $r(X,P)$ be the correlation between the signal currently being broadcast, X, and the pattern part, P , of some production. Each production is assigned a threshold, $T(p)$, that may be changed from time to time. The response of a production to signal X is

$$(3)\qquad \text{Output} = \begin{cases} 0 \text{ if } r(X,P) < T(P) \\ g(r(X,P)-T(P)) \text{ otherwise,} \end{cases}$$

where $g(x)$ is a non-decreasing function of its argument. All the decision mechanism "knows" is the value of the outputs of the various productions. Conflict resolution reduces to determining which production is emitting the strongest output, and allowing that production access to the operations system.

Stated this way, the conflict resolution problem is trivial. It becomes non-trivial when we consider something that is certainly true of all biological systems. All signals are masked by noise. To represent this, noise vectors are added to signals in two places, into the signal broadcast into LTM from the Working Memory and Sensory systems, and into the signal from LTM into the decision system. Obviously this will make pattern recognition more difficult. The size of the correspondence between LTM patterns and working memory will be underestimated, and random fluctuations in matching working memory to irrelevant patterns will be exaggerated. In order to make sure that the appropriate production is selected for action, the decision system must allow for the effects of random variations in the signals that it receives. There are several possible decision policies it can use. All follow the same general procedure; select a signal if it is either momentarily the strongest signal by some margin, M, or if it has maintained itself in position as the strongest signal for some time period, t. The parameters M and t should be set to reflect the relative values of signal and noise in the output of LTM. As a general rule, the greater the noise, the longer it will take to achieve reliable conflict resolution when the signals from LTM are constant. If the signals are variable over time, as they would be if the environment were transient, the amount of noise in the system would determine the maximum speed with which the system could react to environmental change.

A rule for determining T(P), the threshold associated with a production, must be given. Each production in LTM is assumed to be tied to every other production in an associative network (Anderson and Bower, 1973, and many others.) Formally, consider any two productions, s and r, which will be referred to as the sending and receiving productions. There will be some real valued link, a(s,r) that reflects the directional association between them. Whenever the sending production produces some non-zero output, o(s), the receiving production will have its threshold lowered by an amount that depends upon the product, a(s,r) x o(s).

Computationally, what the associative network does is to establish production systems, since the activation of a sending production can sensitize closely associated receiving productions, perhaps to the point that almost any signal will be sufficient to activate them. This is shown in Figure 6. Linkage of productions through association proceeds independently of linkage that may be established through working memory, as the "spread of activation" of a signal from a sending to a receiving production is not dependent upon the sending production's being selected during the conflict resolution stage. This is also shown in Figure 6. This provides two different ways in which one production may activate another, automatic linkage through the association network, and controlled linkage through working memory. The terms 'automatic' and 'controlled' are justified in the sense that controlled activation is directed by the decision system. An analogy to human performance will be considered in the next section.

FIGURE 6 HERE

As a way of summarizing these ideas, the steps in activation of a production system will be traced. The system begins in the state shown in Figure 7, with production P1 linked to P2, P2 to P3 and P4, P3 linked tightly to P5 and loosely to P6, and P4 tightly to P6 and loosely to P5. Further links are also shown. Production P1 matches stimulus pattern S1 in working memory, and the associated action A1 is taken, producing S2 in working memory. At the same time production P2's threshold is lowered to such an extent that action A2 is taken, regardless of the result of the first action, producing S3. Activation through the association network does not discriminate between P3 and P4, hence the determination of the appropriate production will depend upon the extent to which these two productions can match S3. As there is no bias toward the one production or the other, the decision must be made at the conflict resolution stage, and thus will be affected by the efficiency of the decision system. The choice that is made at this point will bias the next decision, between P5 and P6, but will not determine it absolutely, as both productions will have been activated, but in unequal amounts. Thus the last step will be guided both by the "automated" mechanism acting within LTM and by the "controlled" mechanism involving working memory.

- - - - - - - - -

FIGURE 7 HERE

. . . .

APPLICATIONS

A computational theory of thought should provide a plausible framework for tying together disparate phenomena. In this use of the theory a broad brush picture must be painted. One shows that the general form of the phenomena of interest can be produced by the theory, but one does not go to the detail of actually matching theory to data. Such an evaluation will be called a macroscopic evaluation. The theory can also serve as a set of principles to be followed in designing models for specific experimental situations. The models can be tested against observed data. This is a microscopic evaluation. Both steps are important. The advantages of microscopic evaluation have been stressed over and over, particularly within experimental psychology. I believe that microscopic evaluation should be preceded by macroscopic evaluation. Before going to the rather considerable work of generating and testing models for specific situations, we want to be sure that the models themselves are being generated by a theory that is, itself, acceptable on a priori grounds. The remainder of this paper is an attempt to establish the macroscopic acceptability of the CTM just described.

The psychological phenomena to be discussed fall into three broad categories; observations from experimental psychology, the study of individual differences, and biological psychology. Each of these categories will be represented by two problems that have excited considerable study within their own field. In discussing each problem the basic phenomena will first be described, then an explanation will be offered in terms of the model, and some comments will be made about how a computational theory addresses some of the issues raised by the data.

PHENOMENA FROM EXPERIMENTAL PSYCHOLOGY

AUTOMATED AND CONTROLLED RESPONDING. Tasks vary greatly in the amount of attention that they require. "Attention" here has an explicit meaning. Responding requires attention if the response's characteristics are influenced by the presence of distractors, or if making the response interferes with a concurrent task. Schneider and Shiffrin (1977) have reported a particularly dramatic example of automated responding based on the scanning of visual display. Their paradigm is illustrated in Figure 8. An observer was first shown one or more target characters, followed by a very rapidly presented sequence of displays, each containing either a target character and N-1 distractors, or N distractors without a target character. The task was to report whether or not the target was present. When the same characters were targets on some trials and distractors on others, then the speed and accuracy of detection varied as a function of both the number of targets and number of distractors on that trial. The relationship is maintained over thousands of trials. However, when characters were divided into characters used only as targets and characters used only as distractors, the relationship between performance and number of targets and distractors almost (but not quite) disappeared after extensive practice. The detection response was said to be automated.

FIGURE 8 HERE

Automated responding can also be established by extra-laboratory learning. Most of the experiments showing this involve the highly overlearned responding to visual figures that are required in reading. If you are a skilled native speaker of English you cannot look at the visual figure CAT without reading the word "cat." Indeed, it is difficult to inhibit such overlearned responses even, when it would be to your advantage to do so. This is the basis of the Stroop phenomenon (Stroop, 1935).

The distinction between automated responding and its opposite, controlled responding, has been heavily featured in recent literature on attention and performance. Controlled responding is said to be characteristic of poorly learned tasks, and of tasks that involve a great deal of response incompatibility. Automated responding is said to be charcteristic of highly overlearned skilled performance, even where that performance is itself quite complex. I have even heard an argument that automated responding is at the heart of skilled athletic performance.

In the model proposed here controlled responding is defined as responding that involves the decision mechanism and, thus, is guided by information in working memory. Automated responding is defined by the absence of such involvement. A response is automated if the appropriate production is activated with a minimum of processing during the conflict resolution phase. The theory requires that we distinguish two levels of automaticity. The tightest possible connection between a stimulus and response occurs when the stimulus is connected to the response by a single production. An intermediate level of automation can be reached if a sequence of productions is so tightly linked into an association network that activating the first production in the chain will almost always followed by the activation of the subsequent productions.

What sorts of stimulus-response connections can be automated? Shiffrin and Schneider (1977) appear to take the position that virtually any stimulus can become the signal for an automated response, providing that during training the mapping from stimulus to response is always consistent...i.e the same response is required to a fixed stimulus. Others have been more cautious. Treisman and Gelade (1980) have presented evidence indicating that it is difficult, if not impossible, to automate detection responses unless the stimuli are perceptually primitive ones that can be defined by a peripheral sensory process. For instance, it is possible to develop an automated detection response to a red stimulus in a field of blue distractors, or to an F in a field of X distractors, but one cannot automate the detection of a red F in a field of blue F's and red X's.

Results such as those just cited suggest strongly that automated responding is limited to figures that satisfy the perceptual requirements for 'good Gestalt.' Presumably such stimuli are identified at a relatively peripheral level in the sensory system. If such a limitation does apply to automated responding, some of the more ambitious uses of the concept, such as the anecdotal explanation for performance in tennis and basketball, would not be warranted. Some further data that my colleagues and I have obtained in our own laboratory complicates the picture. In one experiment (Yantis, Hunt, and Wright, Note 1) we found that automated detection responses that had been established to upper case letters transferred almost perfectly to the detection of graphically dissimilar lower case letters, such as G and g. Thus we demonstrated generalization of automation along a semantic dimension rather than a physical one. Another experiment (Shimamura and Hunt, Note 2) demonstrated the importance of semantic response dimensions. In this study native readers of Japanese participated in a Stroop type experiment using both Kana (phonetic) and Kanji (ideographic) text. Greater interference was found using the Kanji script, which was related to the interfering (color) response semantically, even though the naming of colors, the overtly interfering response, proceeded more rapidly in Kana. Taken together, these experiments demonstrated automated responding along semantic dimensions of both stimulus and response.

Such findings are not consistent with the position that automated responding is restricted to perceptually primitive units.

The paradox may disappear if we accept the idea that there are two levels of automated responding; one depending on the establishment of a single production tying stimulus to response, and one depending on the establishment of associative links between productions. The first case would lead to a rapid, brief response characterized by a low threshold, and appropriately tuned patterns at the pattern recognition stage, and the emission of a strong single that would rapidly capture (and release) the decision system. This type of automation, however, would be limited to stimuli and responses that could be described within the limits of a single production. Within a computational theory the question "what sort of stimuli can be automated completely" is interpreted as a question about the principled restrictions that should be set upon the complexity of a production's pattern part. On the other hand, there should be no limit to the complexity of an "automated" response based upon a tightly linked chain of productions. It should be noted, though, that even a tightly linked chain would never reach quite the degree of automation as would a single production. Both working memory and the decision system would be required, albeit briefly, to guide the execution of productions in the "automated" chain, and thus there would be some interference with concurrent mental activity.

MENTAL RESOURCES. Why is it difficult to do two things at once? In a trivial case, two activities may be incompatible because they cannot be executed physically at the same time. Try juggling while playing the piano. Mental tasks may be similarly incompatible because two concurrent activities cannot use the same structure. Competition for space in working memory is typically offered as the explanation for our inability to do such things as rehearsing poetry and the multiplication table at the same time. In some situations, though, it is difficult to explain interference as being due to competition for some mental or physical structure. It is difficult, although not quite impossible, to talk while juggling. To offer a more psychological example, performance on the Raven Progressive Matrix test, a widely used test of general intelligence, can be depressed by having a person balance a light lever with one hand while taking the test (Hunt, 1980). Wickens (1979) has listed a number of other examples.

Kahneman (1973) has maintained that we cannot account for non-structural interference without assuming the existence of a "mental resource", somewhat akin to an economic resource, that is shared by concurrent mental activities. Two tasks will interfere with each other if their total resource demands exceed the available supply of the resource. This idea has been elaborated upon in several subsequent papers (Norman and Bobrow, 1975; Navon and Gopher, 1979; Posner, 1978; Hunt and Lansman, in press.) While the need for the concept of a mental resource is widely acknowledged, the nature of that resource has never been specified. Thus it is difficult to connect models of performance that use the concept of mental resources to models of performance that deal with information processing by different structures. This is the sort of macroscopic issue that should be addressed by a general theory.

In computational theory "concurrent activity" occurs when the executions of productions from two production systems, A and B, are interleaved. In this case the tasks associated with the two systems will be said to be "time shared." If we write Pa(i) for the ith production of system A and Pb(j) for the jth production of system B, the sequence of production execution for time shared tasks might be might be:

Pa1, Pa2, Pb1, Pa3, Pb2,Pb3...

From an observer's point of view actions relevant to each task would be executed simultaneously so long as production execution within each task was rapid enough not to produce recordable delays. Interference could arise in several ways. The presence of stimuli from task B in working memory might influence production selection in task A, by altering the pattern recognition phase, mental operators required for one task might be preempted by another, or the proces of conflict resolution might be slowed by the simultaneous activity in production systems A and B.

Two production systems will be defined as being structurally independent if (a) the pattern parts of all productions within one system are random with respect to the pattern parts contained in the other system and (b) if the action parts of the two systems do not use the same (mental) operators. If two production systems are structurally independent interference due to conflicts in the pattern recognition and production execution phase cannot occur. The two tasks may still interfere if the conflict resolution phase becomes a bottleneck. This will be called resource competition, and "mental resources" will be defined as having access to the decision system. Resource competition can occur if the decision system receives two approximately equally strong signals, from two independent systems, both requesting access to the working memory-mental operation system, even though the two productions "want" quite access to different mental structures. This is shown schematically in Figure 9.

FIGURE 9 HERE

Resource competition can be understood by a simple analogy. Imagine an airline reservation clerk who is faced with two customers, each of whom wants to go to a different place, at a different time. There is no competition between passengers for seats on an aircraft, i.e. no structural competition. The customers do compete for the attention of the airline clerk. In particular, suppose that queue discipline is non-existent, so that the customer who shouts the loudest receives immediate attention. The more customers trying to be served at any one time, and the more vocal they are, the harder it will be for any one customer to get service. Furthermore, and this is an extremely important point, the degree of interruption between customers (and between tasks, in the model) will be determined by the disparity between the customers' voices. If one customer is clearly shouting more loudly than another, then the clerk will immediately service the noisy customer. In terms of the theory, the most activated production will be allowed to execute. Resource competition occurs when all the customers shout at an equal level of loudness, thus forcing the clerk to ask them to shout again, so that the loudest voice can be located.

This explanation of interference between concurrently executed tasks is clearly related to the explanation offered for the distinction between controlled and automated tasks. It follows that the more highly automated tasks A and B are, the less use they will make of the decision system, and thus the easier it will be to do them concurrently. On the other hand, the theory states that even the most automated tasks make some use of the decision system. Thus if one kept adding more and more automated tasks into a concurrent activity, there would be a point at which interference would be produced. It might be impossible to do tasks A,B, and C without cost, even though any pair of them could be done concurrently without their interfering with each other.

Two individual differences issues that are the topic of current study by experimental psychologists can be addressed using the concept of attention as a resource. Are some people better than others at doing concurrent activities? In computer science jargon, is there a dimension of individual differences that could be called "time sharing ability"? This apparently simple question turns out to have a complex answer. According to the theory a person who was an efficient time sharer would be a person who executed conflict resolutions quickly, so that the execution of productions from different systems could be interleaved. It is not at all clear, though, that this would show up as a "time sharing ability" in a conventional correlational study of performance. Two activities would most easily be time shared if they were highly automated. In such a case time sharing capability would be related to the ability to do the tasks singly, since degree of automation would predict performance both in the single activity and concurrent activity situations. Suppose that the activities were both controlled tasks. In such a case the tasks, when executed singly, would reflect the efficiency of the decision system, because it is so involved with the execution of controlled tasks. These observations pose a problem for the identification of individual differences in time sharing, because one would always expect performance in single task conditions to be related to dual task performance. This is true even though the postulated mechanism for the control of concurrent activities is a relatively simple one.

Psychologists and educators have long been interested in the difference between expert and novice performance. Larkin et al. (1980) propose that the distinction be modeled by differences in the production systems required to model each type of behavior. It appears that experts can be better modeled by systems that are more tightly linked, so that the detection of a single key characteristic of a problem leads to the immediate execution of a stage of relevant actions, as a unit. To illustrate, once an expert chess player has decided upon an opening the appropriate sequence of moves follows quickly, up to the next choice point. A novice must discover each move.

If this view of expert behavior is combined with with the theoretical explanation of resource competition just offered a prediction follows that, insofar as I know, has not been tested. Figure 10 depicts the production systems of an expert and a novice, doing some abstract task. Double lines have been used to indicate tight linkage. When the expert detects a signal sufficient to activate production P1 productions P2 and P3 will follow automatically. During this period the expert's decision system would be free to respond to extraneous signal, (Px in the diagram) because it would be easier to interleave response to the signal with the execution of automated task relevant activity. The situation changes after the execution of production P3. At this point a choice point occurs, and the expert's decision system will be fully occupied with the discrimination between two strongly activated signals, from P4 and P5. The extraneous signal will not be able to break into the conflict resolution sequence unless it is stronger than either of these. In the novice's case, the discrimination will be between two weak signals, and hence it will be easier for an irrelevant signal to capture the decision system.

FIGURE 10 HERE

Suppose that two people, an expert and a novice, are asked to do a task and, simultaneously, to respond to a probe task, such as turning of a light if it comes on. If the probe is presented at a time at which performance on the main task is routine the expert should respond to the probe task more rapidly than the novice, other things being equal. If the probe is presented at a time at which alternative behaviors are possible in the main task, the expert should be slower to respond to the probe than the novice.

INDIVIDUAL DIFFERENCES ISSUES

THE SOURCE OF INDIVIDUAL DIFFERENCES. To what extent are individual differences in cognitive performance due to physical differences between people, and to what extent are they due to differences in the knowledge that people possess? This question has perhaps engendered more controversy than any other issue in psychology. There is a natural way to frame the question within a computational theory. What individual differences in performance should be mimicked by changing the production systems available to a simulation, and what differences should be mimicked by changing the parameters of the underlying machine?

It is certainly true that all performance depends upon the possession of the appropriate productions by the performing system. Simply having the productions in LTM, however, is not enough. They must be tied together into appropriate production systems. Loosely, the presence of a production only ensures that the system knows how to make a particular reaction to a stimulus. The linkages between productions determine the contexts in which this reaction can occur. It has been amply demonstrated that these two different forms of knowledge represent different dimensions of individual performance in humans. One of the striking findings of recent research on mental retardation, for instance, is that retardates often are capable of executing basic mental tasks, such as rehearsing a list of numbers in order to memorize it, but that they do not recognize the context in which such strategies are appropriate (Robinson & Robinson, 1976, pg.299).

Individual differences in the possession and organization of productions can be thought of as differences in knowledge possession. Individual differences in the effectiveness of production execution at the machine level, differences in what has been called the "mechanics of thought" (Hunt and Poltrock, 1974) are also important. The mechanics vs. knowledge distinction offers a way to think about the classic nature-nurture distinction. Nurture could affect mechanics as well as knowledge, but nature (genetic influences) can act only through alteration of the biological machinery that influences our thinking.

If the computational viewpoint were to be accepted by behavior geneticists it would lead to the design of studies of the "inheritance of intelligence" that are rather different from most of those conducted to date. Most behavior genetics studies use (often elegant) mathematical analyses of genetic models to partition the variance in individual performance on quite complex mental tests, such as the Wechsler scales or the Raven Matrix test discussed earlier. According to a computational theory of thought, it would be more sensible to use as one's dependent variable extremely simple tests that, insofar as possible, tapped a single component of the cognitive computational process. The obvious candidates are tests of the pattern recognition and conflict resolution processes, tests of the efficiency of working memory, and tests of specific mental operations. The knowledge gathered from such studies would be harder to fit into such global questions as "what is the heritability coefficient of intelligence?", but the computational viewpoint does not regard this as a terribly relevant question anyway. It makes more sense to examine individual differences in specific components of mental action.

EXPLAINING THE DIMENSIONS OF INTELLIGENCE. Psychometric theorists break intelligence into a number of components. No attempt will be made to review the various proposed dimensions of intelligence. The interested reader is referred to Nunnally (1978) and Sternberg (1977) for discussions from different points of view. For our purposes, a particularly interesting theoretical distinction is Cattell's (1971) distinction between crystallized (Gc) and fluid (Gf) intelligence. Horn has elaborated upon this distinction in great detail, and the remarks in this section should be understood as referring to Gc-Gf theory as presented by him (Horn and Donaldson, 1979).

Crystallized intelligence (Gc) is assumed to be a dimension of individual performance associated with the ability to use highly overlearned, culturally acquired information in an appropriate fashion. A vocabulary test is the epitome of a task requiring Gc. Indeed, most tasks that are markers of Gc are also verbal tasks, which does present a problem of interpretation. Fluid intelligence is a dimension of performance associated with the ability to solve new problems rapidly and efficiently, pernaps by inventing a solution strategy on the spot. Thus fluid intelligence places more emphasis on the ability to rearrange information, while crystallized intelligence places more emphasis on the ability to retrieve it. Fluid intelligence is often tested by non-verbal intelligence tasks, such as the Raven Matrix test mentioned earlier. In fact, this test is sometimes considered a marker for Gf. Verbal tests of Gf can also be constructed. Such test emphasize verbal reasoning about presented information rather than the retrieval of verbally coded information.

From a computational viewpoint the Gc-Gf distinction is closely related to the distinction between controlled and automated tasks. Performance on tasks requiring Gc should demand (a) that a person have the necessary production systems present and (b) that the situation for their activation be recognized. Once the appropriateness of the production system is recognized, its execution, being based on highly overlearned information, should be relatively automatic and efficient. Tasks demanding Gf are tasks that, by definition, cannot be solved by tightly sequenced systems of productions. Hence a solution must be developed one step at a time. Such tasks should be more influenced by the efficiency of the decision system and of working memory as these systems, taken together, determine the selection of productions in situations in which tightly linked sequences of productions do not exist.

Over the past several years a number of studies have been reported that bear directly on this issue. Tasks that require the simple retrieval of information from long term memory (i.e. tasks that test the pattern recognition system) typically show reliable but moderate correlations (about .3) with psychometric measures of crystallized intelligence. This was first shown for the stimulus matching task (Posner and Mitchell, 1967), in which a person is shown two different forms of the same lexical item...e.g. the pair of symbols CAT and cat...and asked if they have the same name. People who do this rapidly tend to have high scores on Gc tests (Hunt,1978b). More recently we have obtained similar results using other tests of the retrieval of overlearned information, such as the semantic verification (Collins and Quillian, 1970) and lexical identification (Meyer,Schvaneveldt, and Ruddy, 1974) tasks (Hunt, Davidson, and Lansman, in press).

Rather different results are obtained if one uses as the target variable a test of fluid intelligence, such as the Raven Matrix test. While there is much less information available, it appears that the information processing paradigms that are most highly correlated with Raven Matrix performance are tests of choice behavior, as exmplified by the slope variable in a Fitts' type choice reaction study, and tests of rapid verbal reasoning, such as the sentence verification task developed by Clark and Chase (1972). The latter task is believed to have a heavy verbal component. Some typical measures obtained in the study of information processing correlates of Gc and Gf tasks are shown in Table 1. They suggest, but certainly do not prove, that Gf is more closely related to decision making than to pattern recognition.

TABLE 1 HERE

ISSUES IN BIOLOGICAL PSYCHOLOGY

AGING. Botwinick (1977) has summarized the gross observations about age and intellectual change by referring to a "classic aging pattern". As people move through the adult years verbal performance is maintained and even increased, while performance on non-verbal intelligence tests decreases. The picture is not quite that simple, because verbal tests that involve reasoning and inference appear to show a decrease, while tests of simple retrieval and recognition of verbal meanings increases (Cohen, 1979). Horn, in his extensive writings on this subject (see, in particular, Horn, 1978, 1979) has argued that the correct distinction is between Gc and Gf functions rather than between verbal and non-verbal functions. Hasher and Zacks (1979), in an article that is quite compatible with the theory proposed here, claim that as one ages automated functions are retained but controlled processing becomes more difficult.

Why should this pattern of change occur? From the viewpoint of a computational theory of thinking, what functional changes would lead to these particular changes in performance, and how might these functional changes come about? A number of authors have suggested that what may happen is that the aging brain simply is not a reliable transmitter of signals between its various parts. This could occur because the death of nerve cells would, over one's lifetime, produce a progressively less reliable signal, even though the locus of nerve cell death was uncorrelated with the locus of information storage in the brain. Another possibility is that demyelinization of the nerve cells may produce unreliable signal transmission. In terms of a computational model, this would be reflected by an increased error component in the signals transmitted through the sensory, working memory, and decision systems. This would lead either to more erroneous processing if the speed time of responding were to be maintained, or to slower mental processing if the same level of accuracy were to be maintained. It appears more consistent with the literature to assume that the latter is the case, mental processing in the elderly appears to be a (not entirely successful) effort to sacrifice speed of processing in order to retain accuracy (Welford, 1977). Deterioration of performance should be less evident in automated (and Gc) tasks, since unreliability would affect only the pattern recognition stage, whereas controlled (and Gf) tasks would be affected at both the pattern recognition and decision stages.

NEUROPSYCHOLOGICAL PHENOMENA. The study of behavioral changes associated with brain injury provides us with a fascinating glimpse of the specificity of mental functioning. Surprisingly, there has been remarkably little integration of neuropsychological observation into the theoizing of either experimental or psychometric psychology. Possibly this is true because the "average" head injury case has little to tell us beyond reaffirming some general principles, such as the principle that most language functions are located in the left hemisphere. We learn much more from the occasional cases in which small, discretely localized insults to he brain produce highly specific behavioral defects. Examples of such cases are the "deep dyslexic" cases reported by Morton and Patterson (1979), in which patients have selectively lost the ability to deal with English function words, or cases in which selective loss of either long term or short term memory function appears (Shallice, 1979).

The careful study of such cases is of interest in the construction of a computational theory of the mind because neuropsychology gives us some indication of what the basic mental operators are. We know, for instance, that there is a specific operator for storing information in LTM because we know that it is possible to disable this operation, on an amazingly selective basis, by damage to the hippocampus (Milner, 1970). Systematic analysis of selected cases in neuropsychology may enable us to compile an extensive catalog of operators that are "basic" in the sense that they can be selectively physically disabled. Such a catalog would be useful because it would provide us with a vocabulary for stating the action part of productions. This would be an important, and psychologically principled, restriction on the power of production languages.

The view that neuropsychological defects are equivalent to a loss of specific mental operators provides us with an important and in some ways limiting theoretical interpretation of the aftereffects of brain injury. Loss of previously demonstratable functions would be associated with an inability to execute productions, but not to a loss of the productions themselves. Thus loss of function should be largely, if not entirely, a loss of the "controlled" portion of previous learning, i.e. the ability to use a production to produce changes in working memory. Effects associated with the spread of activation from an aroused (but now unexecutable) production to other productions would remain. This is the way in which a computational theory can account for such phenomena as the observation that some neuropsychology patients may have an emotional response to an object that they cannot recognize, and that even when an object is missnamed, the erroneous response is connected to the correct response by some simple semantic rule, such as saying "uncle" for "nephew" (Marshall & Newcombe, 1966).

CONCLUDING REMARKS

In this paper the nature of a computational theory has been emphasized. Such a theory provides a framework within which models of specific tasks can be generated. In general, there will be more than one possible model for a given task. Development of the theory will involve the generation of models that can apply to concrete tasks, and the selection amongst them of models that can describe the data of specific experiments.

If the theory is to be developed further we need models of four key functions; the pattern recognition process, the conflict resolution process, the manner of representation of information in working memory, and of the action of permissible mental operators on this representation. Empirical work to test these models will probably follow the traditions of experimental psychology. We need to find situations in which we can assume that the production system programs to do the task are trivially straightforward, so that we can be sure that behavior is being controlled at the "mechanics level", by the design of the machine executing the program rather than by the design of the program being executed. The emphasis here is on the use of Cognitive Science concepts to explain behavior in the psychology laboratory.

As this work is going on, other cognitive science specialists will be working in quite another tradition. Attempts are being made now, and will continue to be made, to write production system programs that are capable of mimicking the broad brush of human behavior in a variety of very complex tasks. The work of Riesbeck (1975) on the understanding of text is a good example. Hopefully these two lines of study will coalesce. Eventually we would like to see complicated programs for the management of knowledge being executed on (virtual) machines that had been established as reasonable models of the mechanics of human thought. Such a hope can be looked upon as a statement of the ancient goal of reductionism in science. Sociology should be reduced to psychology, psychology to physiology, physiology to chemistry, and chemistry to physics. More realistically, the goal here is to reduce "the psychology of problem solving" to "the psychology of information processing.", and to point toward a connection between information processing and physiological psychology. In the past efforts toward this goal have not been too successful. The reason that the Cognitive Science effort may succeed is that it has a new and powerful set of concepts. The distinction between program, programming language, and machine design may provide the ideas we need in order to establish a reductionist theory of mental action.

REFERENCES

Anderson, J.R. Language, Memory, and Thought. Hillsdale, New Jersey: Erlbaum, 1976.

Anderson, J.R. & Bower, G.H. Human associative memory. New York: Wiley, 1973.

Anderson, J.R., Kline, P. and Beasley, C.M. Jr. A general learning theory and its application to schema abstraction. Carnegie-Mellon University, Department of Psychology Technical Report, 1978.

Botwinick, J. Intellectual abilities. In J. Birren & K.W. Schaie (Eds.) Handbook of the psychology of aging. New York: Van Nostrand Reinhold, 1977.

Cattell, R. B. Abilities: Their structure, growth, and action. Boston: Houghton-Mifflin, 1971.

Clark, H.H. and Chase, W.G. On the process of comparing sentences against pictures. Cognitive Psychology, 1972, 3, 472-517.

Cohen, G. Language comprehension in old age. Cognitive Psychology, 1979, 11, 412-429.

Collins, A. and Quillian, M.R. Does category size affect categorization time. Journal of Verbal Learning and Verbal Behavior, 1970, 9, 432-438.

Hasher, L. and Zacks, R. Automatic and effortful processes in memory. Journal of Experimental Psychology: General, 1979, 108, 356-388.

Horn, J.L. The nature and development of intellectual abilities. In R.T. Osborne, C.E. Noble, & N. Weyl (Eds.) Human Variation: The biopsy of age, race, and sex. New York: Academic Press, 1978.

Horn, J.L. Intelligence-Why it grows, why it declines. In L. Willerman & R.G. Turner (Eds.) Readings about Individual and Group Differences. San Francisco: W.H. Freeman and Co., 1979.

Horn, J.L. and Donaldson, G. Cognitive development II: Adulthood development of human abilities. In O.B. Brim, Jr. and J. Kagan (Eds.) Constancy and change in human development. New York: Academic Press, 1979.

Hunt, E.B. A model of information acquisition and organization. Symposium on Human Information Processing. XVIII International Congress of Psychology, Moscow, 1966.

Hunt, E.B. What kind of computer is man? Cognitive Psychology, 1971, 2, 57-98.

Hunt, E.B. Imageful thought. In J. Cotton and R. Klatzky. (Eds.) Semantic factors in cognition. Hillsdale, N.J. Erlbaum Associates, 1978.

Hunt, E.B. The design of a robot mind: A theoretical approach to issues in individual differences. In M. Freidman, J. Das, and P. O'Connors. Intelligence and learning, London, Plenum, 1981.

Hunt, E.B. The mechanics of verbal comprehension. Psychological Review, 1978, 85, 109-130.

Hunt, E.B. Intelligence as an information processing concept. British Journal of Psychology, 1980, 71, 449-474.

Hunt, E.B., Davidson, J., & Lansman, M. Individual differences in long-term memory access. Memory and Cognition., in press.

Hunt, E.B. and Lansman, M. Individual Differences in Attention. In R. Sternberg (Ed.) Advances in the Psychology of Human Intelligence, Vol. 1, 1981.

Hunt, E.B. and Poltrock, S. The mechanics of thought. In B. Kantowitz (Ed.) Human Information Processing: Tutorials in Performance and Cognition. Potomac, Maryland: Erlbaum Associates, 1974.

Johnson-Laird, P.N. Mental models in cognitive science. Cognitive Science, 1980, 4, 71-114.

Kahneman, D. Attention and Effort. New York: Academic Press, 1973.

Luria, A.R. Higher cortical functions in man. London. Taulstuck, 1966.

Marshall, T.C. and Newcombe, F. Syntactic and semantic errors in paralexia. Neuropsychologia, 1966, 4, 169-176.

McDermott, J. and Forgy, C. Production system conflict resolution strategies. In D.A. Waterman and F. Hayes-Roth (Eds.) Pattern directed inference systems. New York: Academic Press, 1978.

Meyer, D.E., Schvaneveldt, R.W. and Ruddy, M.G. Loci of contextual effects on visual word recognition. In P.A. Rabbitt and S. Dornic. Attention and Performance V. London: Academic Press, 1974.

Miller, G.A. The magic number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.

Milner, B. Memory and the medial temporal regions of the brain. In K. Pribram and D.E. Broadben, (Eds.) Biology of Memory. New York: Academic Press, 1970.

Morton, J. and Patterson, K. Little-words, No! In Coltheart, M., Patterson, K.W., and Marshall, J.C. (Eds.) Deep Dyslexia. London, Routledge and Kegan Paul, 1980.

Navon, D. and Gopher, D. On the economy of the human processing system. Psychological Review, 1979, 86, 214-255.

Newell, A. Physical symbol systems. Cognitive Science, 1980, 4, 135-184.

Newell, A. Artificial Intelligence and the concept of the mind. In R. Schank and K. Colby (Eds.) Computer Models of Thought and Language. San Francisco: Freeman, 1973.

Newell, A. and Simon, H. GPS, A program that simulates human thought. In E. Feigenbaum and J. Feldman (Eds.) Computers and Thought. New York: McGraw-Hill, 1963.

Newell, A. and Simon H. Human Problem Solving. Englewood Cliffs, New Jersey. Prentice-Hall, Inc. 1972.

Norman, D.A. and Bobrow, D.B. On data limited and resource limited processes. Cognitive Psychology, 1975, 7, 44-64.

Nunnally, J. Psychometric theory (2nd Ed.) New York: McGraw-Hill, 1978.

Posner, M.I. Chronometric explorations of mind. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1978.

Posner, M.I. and Mitchell, R. Chronometric analysis of classification. Psychological Review, 1967, 74, 392-409.

Pylyshyn, Z. Computation and cognition issues in the foundations of cognitive science. The Behavioral and Brain Sciences, 1980, 3, 111-169.

Quinlan, J.R. and Hunt, E.B. A formal deductive problem-solving system. Journal of Associate Computing Machinery, 1968, Vol. 15, No. 4, 625-646.

Riesbeck, C. Conceptual Analysis. In P.C. Schank (Ed.) Conceptual Information Processing. Amsterdam: North Holland Publishing, 1975.

Robinson, N. and Robinson, H. The Mentally Retarded Child. (2nd ed.) New York: McGraw-Hill, 1976.

Schneider, W. and Shiffrin, R. Controlled and automated human information processing I. Detection, Search, and Attention. Psychological Review, 1977, 84, 1-66.

Shallice, T. Neuropsychological research and the fractionation of memory systems. In L.G. Nilsson (Ed.). Perspectives in Memory Research. L. Erlbaum Associates: Hillsdale, N.J., 1979.

Shimamura, A. and Hunt, E.B. Differences in Processing Ideographic (Kanji) and Phonemic (Kana) Symbols: A Japanese version of the Stroop Test. Paper presented at the Western Psychological Assn. Meeting, May, 1980, Honolulu, Hawaii.

Sternberg, R.J. Component processes in analogical reasoning. Psychological Review, 1977, 84, 353-378.

Stroop, J.R. Studies of interference in serial verbal reactions. Journal of Experimental Psychology, 1935, 18, 643-662.

Treisman, A. and Gelade, G. A feature-integration theory of attention. Cognitive Psychology, 1980, 12, 97-136.

Welford, A.T. Motor performance. In J.E. Birren & K.W. Schaie (Eds.). Handbook of the Psychology of Aging. New York: Van Nostrand Reinhold, 1977, 450-496.

Wickens, C.D. The structure of attentional resources. In R.S. Nickerson (Ed.) Attention and Performance VIII. Hillsdale, N.J.: L. Erlbaum Associates, 1979.

Yantis, S., Hunt, E.B., and Taylor, J. Generality of Automatic Information Processing. Technical Report, Department of Psychology, University of Washington, 1980.

PROGRAM - LANGUAGE - MACHINE

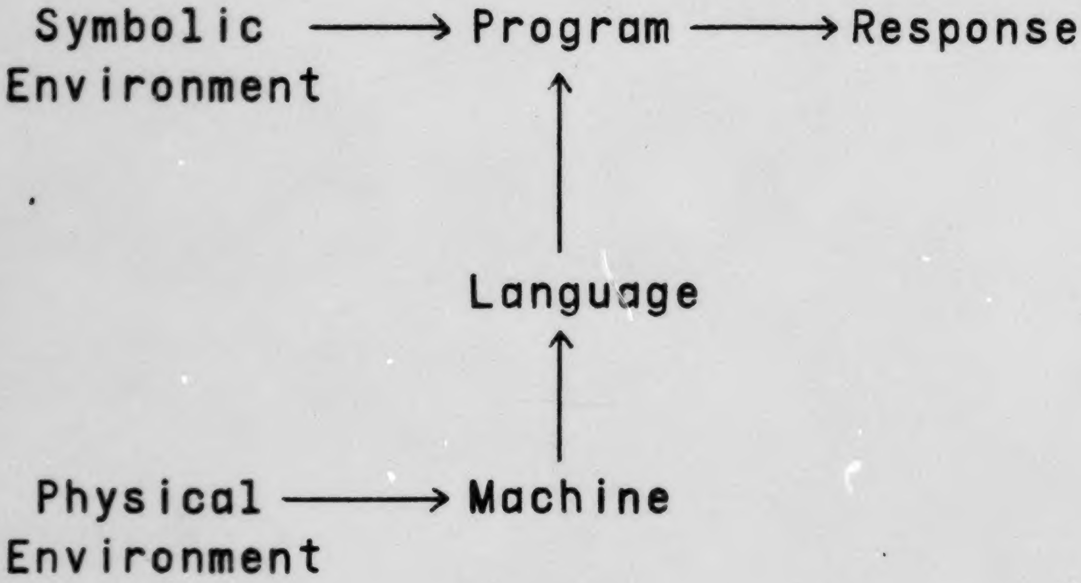

Symbolic Environment → Program → Response

Language → Program

Machine → Language

Physical Environment → Machine

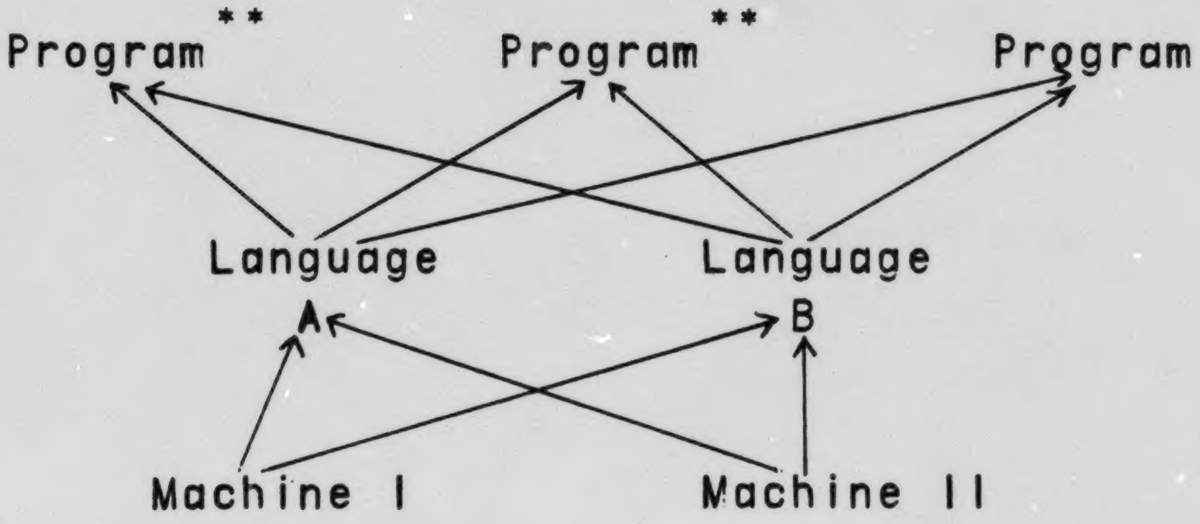
Program**
Program**
Program
Language A
Language B
Machine I
Machine II

THE EXCESS POWER PROBLEM

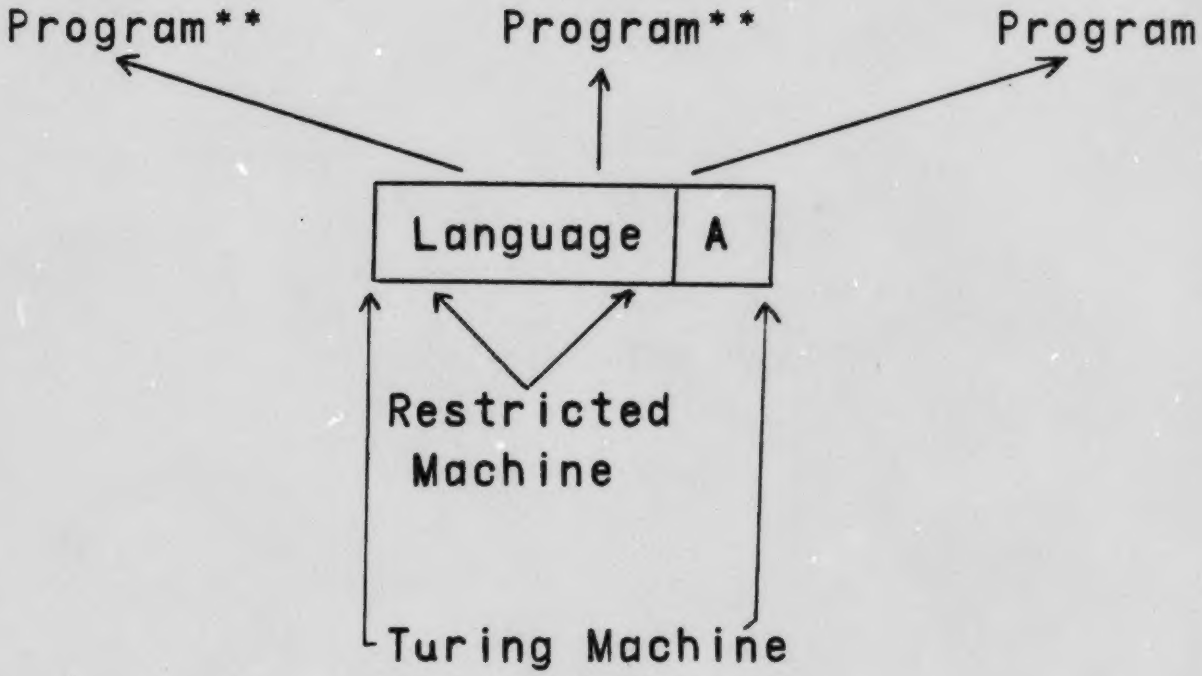

Program**

Program**

Program

Language

A

Restricted Machine

Turing Machine

IMPLEMENTATION RESTRICTIONS

A PRODUCTION CYCLE

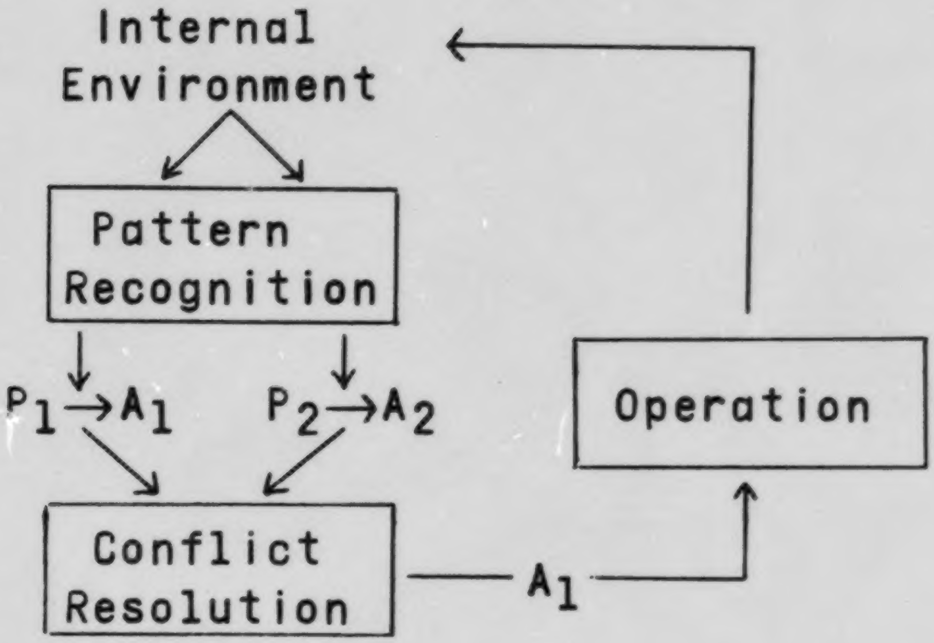
Internal Environment
Pattern Recognition
$P_1 \rightarrow A_1$
$P_2 \rightarrow A_2$
Conflict Resolution
A_1
Operation

103

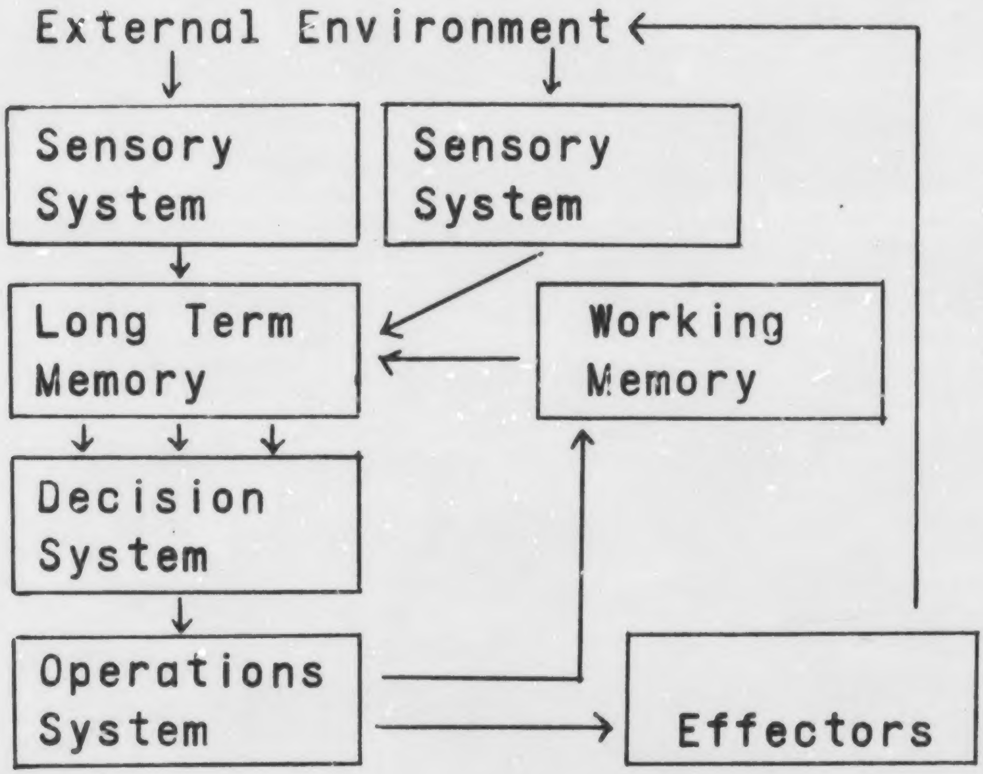
External Environment
Sensory System
Sensory System
Long Term Memory
Working Memory
Decision System
Operations System
Effectors

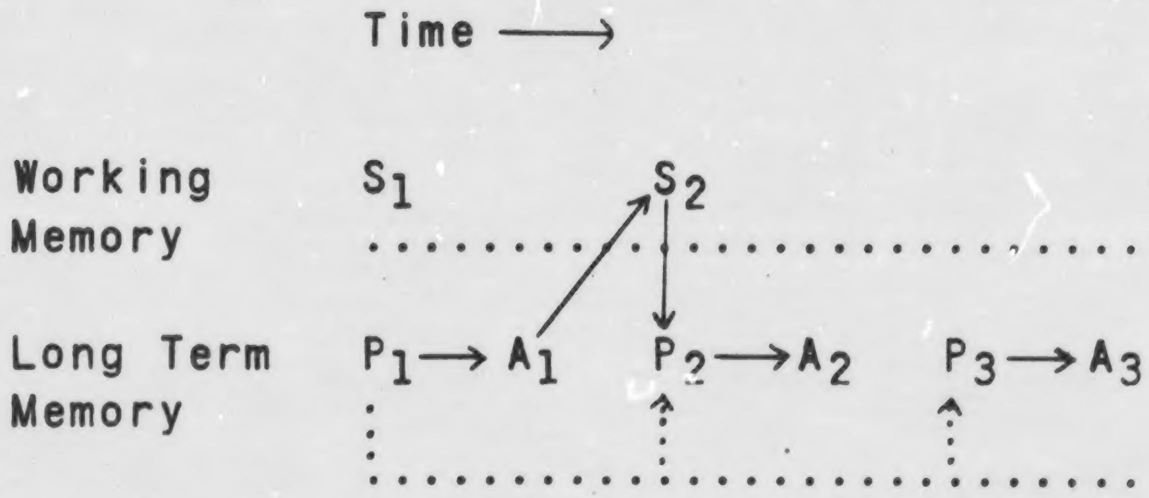
Time →
Working Memory
S1
S2
Long Term Memory
P1 → A1
P2 → A2
P3 → A3

SEQUENCE OF PRODUCTION ACTIVATIONS

PRODUCTIONS LINKED INTO A SYSTEM

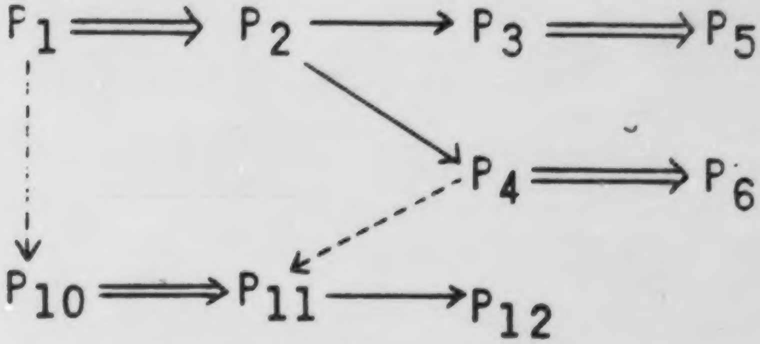
P1
P2
P3
P5
P4
P6
P10
P11
P12

VISUAL DETECTION

Target

X, T

Display

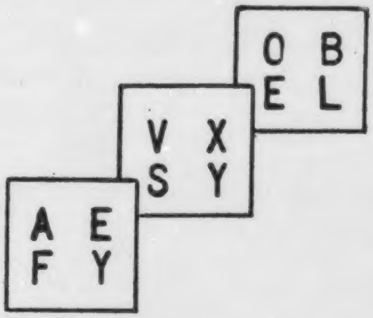
O B
E L
V X
S Y
A E
F Y

RESOURCE COMPETITION

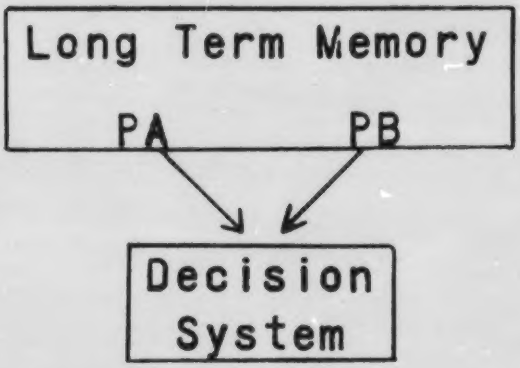

Long Term Memory

PA

PB

Decision System

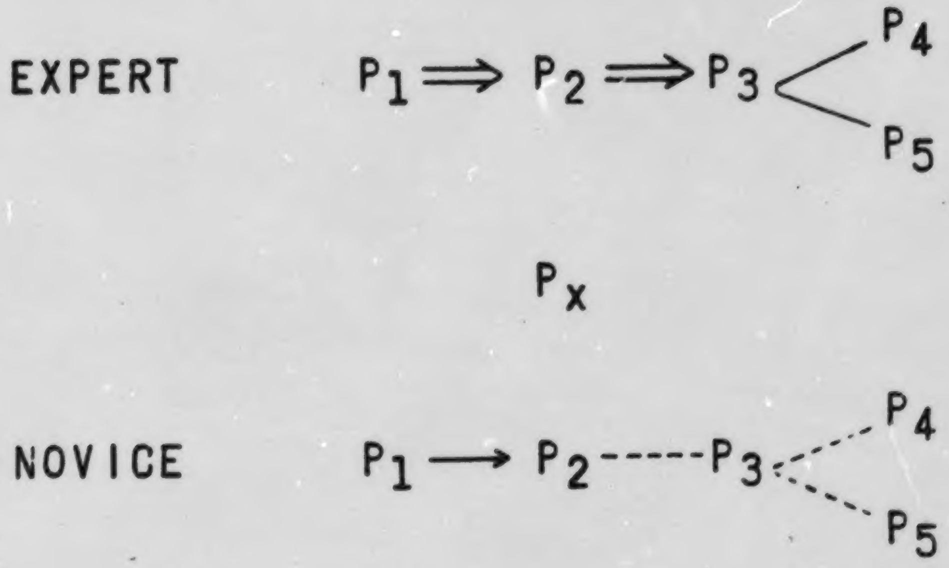
EXPERT
P1
P2
P3
P4
P5
Px
NOVICE
P1
P2
P3
P4
P5

TABLE ONE

	PSYCHOMETRIC VARIABLE	
EXPERIMENTAL TASK	Reading comprehension, Vocabulary	Raven Matrices
Lexical identification Semantic categorization	moderate	small
Speed of choice	??	moderate

Types of Experimental Paradigms that have been related to performance on psychometric tests

BRAIN MECHANISMS IN HUMAN INFORMATION PROCESSING: IMPLICATIONS FOR A COGNITIVE THEORY OF PROCESSING RESOURCES

Jackson Beatty
Department of Psychology and Brain Research Institute
University of California, Los Angeles

Paper first presented at the ONR Lake Wilderness Attention Conference.
Portions of this research were sponsored by the Office of Naval Research under Contract N00014-76-C-0616.

BEST AVAILABLE COPY

1.0 Introduction.

The human nervous system, like any physical information processing system, is limited in capacity. Capacity limitations are unimportant when processing demands are low, as all demands can be easily met. However, people are frequently required to process more information at higher rates than their nervous systems can accommodate, a situation that leads either to a restriction in the tasks that are attempted or to a necessary deterioration in performance quality. It is not surprising that the nature of information processing limitations has formed a central focus for cognitive psychology in the past quarter century.

A number of views of the nature of processing limits have been proposed by cognitive psychologists. Broadbent, in his seminal book Perception and Communication (1958) saw the critical limitation in processing as a single, non-sharable information channel linked to the senses by a selective filter and passing its output to memory and response systems. However, evidence of concurrent, high-level processing of semantic information from multiple speakers led Moray (1967) to proposed that capacity may be shared between different processes or tasks. Kahneman (1973) adapted Moray's general capacity model to an expandable general capacity theory, persuaded by physiological evidence that processing capacity may increase with increasing demands for that capacity. Kahneman identified capacity with attentional processes, a view that is shared by many contemporary theorists.

These ideas were formalized and generalized by Norman and Bobrow

(1975), who introduced the concept of multiple processing resources in their brief treatment of time-sharing behavior:

> "Any information processing device has programs and some mechanism for executing those programs. When a program is executed, it requires input data and it consumes resources. A set of programs that is being executed for a common purpose and for which resources are allocated as a unit is called a process. Resources are such things as processing effort, the various forms of memory capacity, and communications channels. Resources are always limited (p. 45)."

The lack of convincing behavioral evidence in support of a single processing resource has encouraged the development of multiple resource theories (Navon and Gopher, 1979, 1980; Sanders, 1979; Wickens, 1979, 1980). However, the persistent difficulty with all multiple resource theories is the identification and specification of putative processing resources.

It would seem, after several decades of research on the problem of processing limits that there would be some concensus in the field regarding the most appropriate type of model, but this is not the case. Sanders (1979) in his recent, thoughtful review found compelling evidence in support of the single channel hypothesis, Kahneman's effort model and multiple resource theory; none of these models appears to be completely satisfactory.

Three fundamental questions appear to confront any theory of processing resources:

1. Is there a single general processing resource or are processing resources multiple?

2. How can a processing resource be identified and specified?

3. Is the processing capacity of a resource fixed or partially expandable as Kahneman suggested?

It appears that these issues can be resolved only by empirical evidence, but the behavioral data presently available are insufficient in themselves to provide an acceptable solution.

A complementary approach to these questions is that of human neurophysiology. Since the processing resources in human cognition must be functions of the human nervous system, human neurophysiology may provide an important source of theory and data for the study of processing resources. Human neurophysiology, in my opinion, may make three types of contributions to the development of cognitive theory. First, it provides an important source of converging data for the testing of theoretical positions. For example, in later sections of this chapter, evidence concerning each of three questions mentioned above will be reviewed. In this way, neurophysiological data can help resolve some of the more difficult questions in cognitive theory. Second, neurophysiological data may provide important constraints on the types of cognitive theories that are to be considered. Proposing theories of human cognition that are incompatible with what is known about the human nervous system is not good psychology. Finally, neurophysiological data might serve a heuristic role in the development of cognitive theory, suggesting specific hypotheses that are amenable to behavioral testing. A few such suggestions appear later in the chapter. Thus, I believe, human neurophysiology may serve as an important -- and often neglected -- source of information in developing a more satisfactory theory of the structure of human cognitive processing.

Three topics of recent neurophysiological inquiry seem relevant to the study of processing resources in the human brain. The first involves the use of metabolic mapping techniques to study cerebral energy metabolism during information processing in the human nervous system. This new literature suggests that processing resources may be expandable in the sense proposed by Kahneman (1973). The second, related literature involves the measurement of regional cerebral blood flow during human information processing. These data support the idea that there are multiple, specific processing resources and provide heuristic information for their identification. The third line of neurophysiological research involves peripheral indications of reticular core activation during the performance of cognitive tasks. These data support Kahneman's original effort hypothesis and suggests a hierarchical model of the structure of processing resources. Together, these data may serve to clarify the problematic questions of human information processing resources, their nature, their structure and their functions.

2.0 Cerebral Energy Metabolism and Human Information Processing

Although the human brain accounts for only 2 percent of the body's weight, it utilizes 20 percent of the total resting oxygen consumption of the body. If the flow of blood to the brain is interrupted, all available brain oxygen reserves are depleted within 10 seconds and unconsciousness results. Brain oxygen levels are critical because oxidative glucose metabolism is the only source of energy for the neurons of the central nervous system and processing information within the CNS requires energy.

Information transactions within the human brain depend upon the actions of large populations of individual neurons, which communicate among themselves by synaptic interactions. Further, information is communicated between different brain regions by spike potentials generated in the excitable membranes of the cell body and axons. Thus, the information processing functions of the central nervous system may be linked to its energy consumption at several levels (Krnjevic, 1975):

1. The sodium-potassium pump. The standing membrane potential which provides the basis for all information processing at the cellular level is the product of an active, energy consuming membrane process that acts to transport sodium out and potassium into the neuron. In doing so, it is moving paired ions against both concentration and electrical gradients, an action that requires significant amounts of metabolic energy. The generation of an action or spike potential in a neuron is effected by momentarily

permitting the exit of potassium and the influx of sodium, which creates future work for the sodium-potassium pump. Thus, the metabolic load placed upon this membrane pumping system is a direct linear function of the firing rate of each neuron in the central nervous system. This is one way in which function and metabolic load are coupled.

2. Synaptic events. Many synaptic events are active neurochemical processes requiring energy in their function. For example, in synaptic endfeet all available adenosine triphosphate (ATP), the final energy-rich product of cerebral glucose metabolism, is completely utilized and replaced every 3 seconds. Metabolic energy appears to be required for neurotransmitter release, reuptake and deactivation. Similarly, active processes are involved in the response of the post-synaptic membrane to neurotransmitter release.

3. Synthetic processes. Energy is required in the maintenance of proper supplies of neurochemicals utilized in cellular information transactions. These include the manufacture of neurotransmitter substances, neuromodulators and other molecular neural supplies.

These and other, similar relations link cellular information processing functions with cellular metabolic demands. It is for this reason that measures of cerebral oxidative glucose metabolism provide a primary indication of the intensity of information processing within selected regions of the nervous system. The process of glucose metabolism begins with glucose uptake by neurons as a function of their metabolic demands.

This process was once thought to depend primarily upon passive diffusion across the membrane, but it now appears to involve fast acting, active transport channels (Bachelard, 1975). These channels are sensitive to the molecular structure of glucose, transporting some sugars and not others. But most important is the fact that the glucose transport mechanism is regulated with respect to rate and may, in fact, provide a primary mechanism for controlling the cellular metabolic rate for glucose.

The principal metabolic pathway for glucose energy release involves the phosphorylation of glucose to glucose 6-phosphate by the enzyme glucohexokinase. The final result of oxidative glucose is the production of energy-rich ATP. As each molecule of glucose is oxidized, 36 molecules of ATP are produced. It is the ATP that directly provides the source of energy for the conduct of information processing functions in neural tissue.

2.1 2-Deoxyglucose Measurement of Cerebral Glucose Metabolism

Early attempts to measure brain metabolic activity relied upon comparisons of blood gases and energy substrates in the arterial supply and the venous return from the brain. The results were uninformative: functional aspects of central nervous system activity did not appear to affect cerebral metabolic rates. This lack of metabolic responsivity gave rise to the idea that brain metabolism is unrelated to brain function, in much the same way a computer's energy requirements are unrelated to the details of the information transactions taking place.

The problem, of course, stems from the fact that the brain is not a functionally uniform structure, but is intricately differentiated. Different informational transactions within the nervous system involve different portions of the brain. Thus, if some brain structures were increasing and others decreasing their metabolic demands, little effect upon the aggregate arterial-venous differences would be expected.

It was therefore a finding of considerable importance when a method for accurately measuring cerebral metabolic rate for glucose (CMRGlc) with microscopic spatial resolution was reported by Sokolov and his collaborators (See Sokolov, 1979, for a detailled description of this methodology). By using a radioactively labelled analog of glucose that is only initially competitive in the glucose metabolic pathway, Sokolov was able to obtain photographic images of brain tissue in which image density is proportional to the local rate of glucose metabolism.

2-deoxyglucose (2-DG) is an analog of glucose that is completely competitive with glucose for access to both the membrane glucose transport channels and phosphorylation by glucohexokinase. However, the product of the phosphorylation, 2-deoxyglucose-6-phosphate, does not compete with glucose-6-phosphate for further processing. Thus, when 2-DG is injected into the systemic circulation, a metabolically active neuron will take up 2-DG in proportion to its metabolic rate for glucose. But because the metabolism of 2-DG is blocked at an early stage, 2-DG will accumulate in neurons and not be returned to the systemic circulation as metabolic byproducts. Thus the amount of 2-DG present in a neuron is a joint function of the amount of 2-DG in the general circulation, the length of time following injection, and the average CMRGlc for that cell during the period of 2-DG availability.

The 2-DG method for CMRGlc measurement was applied to the study of the human brain by Phelps and his colleagues (Phelps et al., 1979, 1980, 1981). Using a fluorine-labelled 2-deoxyglucose (FDG) human volunteers are intravenously injected and placed in experimental or control conditions for a period of 20-40 minutes to permit accumulation of the labelled compound in metabolically active brain regions. A noninvasive positron emission scanning device, which is similar in many was to the now familiar CAT scanner, is used to measure and locate the sources of radioactive particles. From these data, computer procedures are employed to derive density tomograms, or computed planes of section through the brain, that reflect the local CMRGlc. At present, the resolution of the reconstructed images is approximately 1.5 cubic centimeters. It should be noted that the actual dose of radiation involved is rather small. Furthermore, the effects of radioactivity are short-lived, the half-life of fluorine being approximately 110 minutes.

2.2 Cerebral Glucose Metabolism in the Resting Human Brain

Determination of the basal rates of cerebral glucose metabolism is a necessary first step in the study of modifications of CMRGlc produced in specific brain regions during information processing. Such data have recently been published by Mazziotta et al. (1980) for the human brain at rest.

Mazziotta and his collaborators employed the FDG tracer and positron emission scanning techniques to measure regional CMRGlc in 7 neurologically normal male university students under conditions of minimal sensory and motor stimulation. Figure 1 presents two-dimensional positron scans of the upper body taken 40 minutes after FDG injection.

Insert Figure 1 Here.

As in the tomographic images, the darker portions of the picture represent areas of higher glucose metabolism. From Figure 1, it is apparent that heart and brain are organs with exceptionally high rates of glucose metabolism.

2.3 Visual Information Processing and Cerebral Glucose Metabolism

The close relation between higher brain function and brain metabolism or work is clearly illustrated in recent investigations of the metabolic changes occurring during visual information processing in man and animal. For example, Phelps and his colleagues (Phelps et al., 1980) examined the effects of the complexity of visual stimulation on brain CMRGlc using the FDG tracer and positron emission scanning imagery. The first part of the experiment involved comparing the effects of different levels of visual stimulation in a group of 15 neurologically normal, young, right-handed, male university students. Four conditions were tested in different subsets of this sample: 1) no visual stimulation produced by closing and/or patching the eyes, 2) a brightly illuminated unpatterned visual field, 3) patterned stimulation, a black and white checkerboard pattern phase reversing at the rate of 2/sec, and 4) complex visual stimulation resulting from walking out of doors in a park.

Using the FDG tracer and positron emission scan imaging, a series of tomograms were obtained that included the full extent of visual cortex. In these tomographic sections, major structural features of the human brain can be distinguished. For example, the gray matter of the cerebral cortex may be seen as an area of highly active metabolism, in contrast to the relatively hypometabolic white matter that lies beneath it. From these images, quantitative estimates of CMRGlc were obtained separately for left and right primary (Brodmann's area 17) and association (Brodmann's area 18 and 19) visual cortex.

Visual stimulation significantly increases the rate of glucose utilization in those cortical regions known to be primarily responsible for processing visual information. Furthermore, the demand for glucose increases as a function of processing complexity. For example, Figure 2 presents a representative set of tomograms obtained with the eyes closed and with white light stimulation.

Insert Figure 2 Here.

It may be observed that the only changes in CMRGlc between these two conditions occur in the visual cortex of the occipital lobe. The darkening of these areas under stimulation conditions reflects the increased metabolic demands imposed by information processing in these areas.

A quantitative summary of these results is presented in Figure 3.

Insert Figure 3 Here.

Relative to the eyes-closed baseline condition, simple white light stimulation resulted in an increase in CMRGlc of about 12 percent in primary visual cortex (PVC) compared with about 6 percent in association visual cortex (AVC). This difference probably reflects the known fact that AVC is more concerned with processing spatial patterns than intensity discrimination. The addition of patterned stimulation (the alternating checkerboards) increased CMRGlc in both areas by approximately 25 percent as compared with eyes-closed baseline. When patterned stimuli were presented to only one of the two eyes, there was a 37 percent decrease in metabolic response in PVC and an 18 percent decrease in AVC compared with binocular stimulation. The relative magnitude of these decreases probably reflects the proportion of monocularly driven cells in these two areas; that is, in PVC about one third of all neurons could be expected to be monocularly driven by the unstimulated eye whereas in AVC most neurons have binocular input.

Finally, the complex visual stimulation of the park resulted in very large increases in CMRGlc. In PVC mean CMRGlc increased by 45 percent and in AVC by 59 percent. Local increases within these areas were as large as 100 percent of baseline metabolic rate. These data demonstrate that the metabolic demands of anatomically and functionally defined regions of human visual cortex increase markedly in response to information processing demands. These increases in CMRGlc are not global but are limited to the relevant regions of cortex. Information processing in visual cortex clearly requires brain work.

All of the changes in CMRGlc induced by visual information processing in normal young adults were bilateral. No differences were present between the right and left cerebral cortices. In contrast, Phelps et al. also measured CMRGlc in seven neurological patients with homonymous hemianopsia, a condition in which the visual input to one of the two cerebral hemispheres is absent. This usually results from unilateral damage to the subcortical visual projection system. Figure 3 shows that, when compared with the normally innervated hemisphere, the visually deprived PVC and AVC are hypometabolic. CMRGlc for denervated visual cortex is markedly suppressed, again illustrating the link between neuronal function and neuronal metabolism.

2.4 Implications for a Theory of Processing Resources

These data have several implications for any theory of human information processing resources. First, they demonstrate that selected task-relevant regions of cerebral cortex become activated during specific information processing tasks. Introducing visual stimulation had no effect on CMRGlc in the other classical sensory areas or association cortices for example. This finding suggests that restricted cortical regions serve as processing resources in the sense proposed by Norman and Bobrow (1975). Further, the data imply that cortical processing resources are specialized and multiple, a theme that will form the focus of the following section.

The second fact emerging from the study of cerebral glucose utilization during visual information processing is that cortical resources are not continuously employed, but only become active when a cognitive task demands the services of that region. Thus, cortical resources are unlike computer processors, which usually grind their electronic wheels continuously, whether or not there is any real work to do. In this sense, cortical resources are not fixed in capacity, but rather are expandable, as Kahneman proposed (1973).

This concept is in accord with the analysis of resource utilization and processing costs offered by Navon and Gopher (1979). Navon and Gopher suggested that resources may have a cost associated with their utilization and proposed that such a cost might be the consumption of "mental energy." The present data suggest that the cost metaphor need not be so loosely applied; the utilization of cortical resources does indeed have a significant energy cost, one that is measurable in the rate of cerebral glucose metabolism.

Thus, recent investigations of task-induced changes in CMRGlc provide support for the idea of multiple processing resources, expandable processing resources and real costs associated with resource utilization.

3.0 Regional Cerebral Blood Flow and Cerebral Specialization of Function

A second approach to the metabolic mapping of regional cortical activation in human information processing concentrates on the cerebrovascular effects of the byproducts of neural metabolism. Cerebral blood flow is locally regulated in a manner that assures an adequate rate of oxygen uptake and carbon dioxide removal in the local cerebral microenvironment. Cerebral blood flow is said to be autoregulated, in that it is determined by brain metabolic demands rather than systemic cardiovascular state. Thus, local cerebral blood flow is closely linked to the local rate of oxidative metabolism, which in turn reflects the local level of neuronal information transactions. The primary controlling mechanism for cerebral blood flow is carbon dioxide, an end product of oxidative metabolism. Carbon dioxide appears to exert a direct vasodilatory effect on the cerebral vasculature and the magnitude of this effect is large: increasing blood carbon dioxide by breathing a 5 percent mixture of carbon dioxide increases cerebral blood flow by 50 percent. A 7 percent mixture results in a 100 percent increase in cerebral blood flow. Blood carbon dioxide levels sets the vasodilatory tone and therefore the basal rate of cerebral blood flow.

The effect of local oxidative metabolic rate on blood flow operates against this background. As cerebral information processing increases the local metabolic rate, the output of carbon dioxide in the extracellular fluid of the brain increases producing a change in cerebral pH in the direction of acidosis. It is believed that the increase in extracellular, extravascular hydrogen ions mediates the local vasodilation of the cerebral microvasculature, resulting in a local increase in cerebral blood flow (Ingvar and Lassen, 1976). For this reason, measures of regional cerebral blood flow (RCBF) may be used to assess the regional involvement of cortical tissue in human information processing.

3.1 Measuring Regional Cerebral Blood Flow.

Radioactive labelling methods provide the basis for most measurements of human regional cerebral blood flow, just as they were utilized in measuring CMRGlc. Most measures of RCBF employ radioactively labelled xenon, an inert gas that diffuses freely across the blood brain barrier. In one version of the method, a bolus of labelled xenon is injected into one of the two common carotid arteries supplying the greater part of one cerebral hemisphere. The gas immediately passes from the blood into brain tissue, where its presence is detected by banks of radiation detectors placed against the scalp. The rate of regional cerebral blood flow in the grey matter of the cortex is estimated from the rate at which this labelled gas diffuses back from brain to blood and is carried off and lost in the volume of the body. From the first two minutes of this so-called washout curve, the rate of RCBF is estimated. If RCBF is high, then the larger volume of blood flowing through the region washes out the diffused xenon more rapidly. If RCBF is low, the xenon disappears less quickly. By the use of careful measurement methods, an absolute rate of RCBF may be obtained.

Although RCBF and CMRGlc measures index complementary aspects of cerebral oxidative metabolism linked to brain work, the primary advantage of the RCBF methods is that they are much older and therefore have had the opportunity to provide a substantive body of knowledge concerning regional cerebral involvement in human information processing. It is to this literature that we now briefly turn our attention.

3.2 RCBF in the Resting Brain.

In the absence of any explicit task the pattern of RCBF is far from uniform; flow rates for frontal cortex are 40 to 50 percent greater than in posterior regions. Ingvar (1976) has termed this resting distribution the hyperfrontal pattern and has suggested that it reflects a state of planning on the part of the subject. No real evidence is available that clarifies the nature of cognitive processing during wakeful rest, but there is little question as to the existence of the hyperfrontal pattern. Since the resting distribution is not uniform, it has become common to characterize task-evoked patterns of RCBF as deviations from the resting hyperfrontal pattern (Lassen, Ingvar and Skinhoj, 1978).

3.3 Sensory Processes.

The presentation of simple sensory stimuli in the absence of any explicit cognitive task results in specific increases in RCBF in the appropriate primary sensory cortex and not elsewhere. Lassen, Ingvar and Skinhoj (1978) report that presentation of loud tones resulted in bilateral increases in flow over primary auditory cortex. (See Figure 4 for a mapping of functional cortical areas.)

Insert Figure 4 Here.

Similarly, Lassen, Roland, Larsen, Malamed and Soh (1977) have demonstrated that listening to music also results in bilateral blood flow increases in primary auditory cortex (superior temporal gyrus) and that bilateral increases in occipital RCBF accompanied the viewing of simple visual patterns. Unilateral tactile stimulation elicits contralateral increases in RCBF over the appropriate region of the sensory-motor strip (Lassen et al., 1977). Finally, Foit, Larsen, Hattori, Skinhoj and Lassen (1980) report that stimulation of the median nerve of the arm results not only in the expected increased RCBF of the contralateral sensory-motor hand area, but in the supplementary motor area as well. These may reflect the fact that the median nerve is mixed, containing both sensory and motor fibers.

3.4 Motor Processes.

Several reports have been published demonstrating specific changes in RCBF in motor tasks. Roland, Skinhoj, Larsen and Lassen (1977) have shown that static contraction of the hand produces increased RCBF over the contralateral sensory-motor hand area. In contrast, execution of complex patterns of finger movement induces increased flow both in the contralateral hand area and in the supplementary motor area as well. Imagining the same movements resulted in increases in the supplementary motor area alone, without activation of the contralateral sensory-motor strip. Finally, sequential movement of fingers to target positions elicited contralateral sensory-motor, supplementary motor area and parietal lobe increases in RCBF.

The role of the supplemental motor area in complex movement was further investigated by Orgagozo and Larsen (1979) using RCBF measures. They found that sustained foot contraction produced increases only in the contralateral sensory-motor foot area, whereas complex foot movements resulted in increases in the supplementary motor area as well. Similar patterns were seen for complex finger movements.

Specific regional increases in RCBF have also been reported during eye movements and visual search. Melamed and Larsen (1977) found that eye movements produced bilateral increases in the sensory-motor face area, the frontal eye fields, the supplementary motor area and primary visual cortex.

3.5 Language Processing.

Language processing involves the differential activation of the classical speech areas. These effects often are bilateral. Thus, Lassen, Ingvar and Skinhoj (1978) report that the auditory presentation of simple words activates the superior temporal gyrus (auditory cortex) and Wernicke's area of both cerebral hemispheres. Complex verbal stimuli produce these effects as well as bilateral increases in RCBF over Broca's area in frontal cortex.

The effects of 'non-linguistic' or automatic repetitious speaking were studied by Larsen, Skinhoj and Lassen (1978) who tested 18 right-handed persons while either counting repeatedly to 20 or reciting the days of the week. Focal increases occurred bilaterally in the sensory-motor face and mouth area, supplemental motor cortex, auditory cortex and Wernicke's area. No increase in flow over Broca's area was reported. Perhaps most unexpectedly, there was a relative increase in mean cerebral blood flow over the right, not the left, cerebral hemisphere.

3.6 Implications for a Theory of Processing Resources.

The finding that functionally specialized regions of human cerebral cortex become differentially activated in different types of information processing tasks provides strong support for multiple resource theories. At its highest levels, the nervous system does not function as a generalized pool of processing capacity, although exactly that hypothesis was put forward by Lashley (1929) in his theory of cortical mass action. Instead, different regions within the cerebral hemispheres are specialized to perform qualitatively different information processing functions. The experimental mapping of regional cerebral blood flow reviewed above provides clear evidence of this differentiation in cortical function. Thus, the RCBF literature serves as a source of converging neurobiological data in support of multiple resource theory.

Perhaps more importantly, these data may serve a heuristic function for cognitive theory. The problem of resource identification has plagued the development of multiple resource theories, forcing Navon and Gopher (1979), for example, into a purely theoretical account that ignored the question of resource specification. The literature on regional cerebral blood flow provides a potential solution to the problem of resource identification. These data suggest that each of the three major sensory projection areas with their adjacent association areas may be differentially activated during information processing. Thus, they may constitute putative processing resources. Similarly, the motor strip and adjacent premotor cortex may also constitute a cortical resource, but the distinction between motor and somatosensory cortex is far from clear. Further, the supplementary motor area of prefrontal cortex appears to serve as a coordinating or executive region for the execution of complex, voluntary patterned movements. Finally, the interconnected Wernicke and Broca's areas possibly in conjunction with homologous regions in the non-dominant hemisphere are differentially activated during language processing. These distinctions are important for cognitive psychology as they provide physiological evidence relevant to the problem of resource identification. If two cognitive tasks involve the same cortical regions, they may be thought to invoke the same combination of cortical processing resources. By studying the similarities in tasks which activate the same cortical regions, major advances may be made in transforming multiple resource theories from formal constructions that lack substantive content into a powerful and specific theory of the structure of human cognition.

4.0 Reticular Activating System and Cortical Processing Resources.

Considerable attention has been paid to the question of the involvement of subcortical structures in the complex information processing functions of the human brain. This line of research stems from Moruzzi and Magoun's (1949) now classical discovery of the functional properties of the brainstem reticular formation. The interaction of cortical and subcortical structures in higher cognitive functions has been elegantly summarized by Luria, the late Russian neuropsychologist:

> "The maintenance of the optimal level of cortical tone is essential for the organized course of mental activity. (However,) the structures maintaining and regulating cortical tone do not lie in the cortex itself, but below it, in the subcortex and brain stem. (These structures together are the reticular formation.) Some of the fibres of this reticular formation run upwards to terminate in higher nervous structures such as the thalamus, caudate body, archicortex and, finally, the structures of the

> neocortex. (They play) a decisive role in activating the cortex and regulating the state of its activity. Other fibres of the reticular formation run in the opposite direction. (These descending fibers) subordinate the lower (brainstem) structures to the control of programmes arising in the cortex. The higher levels of the cortex, participating directly in the formation of intentions and plans, recruit the lower systems of the reticular formation of the thalamus and brain stem, thereby modulating their work and making possible the most complex forms of conscious activity." (from Luria, 1973, pp. 45-60)

The dynamic interaction of cortical structures and the reticular core as an essential feature of the neurophysiology of cognitive processing is also emphasized by Brodal (1981) in his recent authoritative review of the neuroanatomy of the reticular activating system:

> "The cerebral influence on the reticular formation can hardly be overrated. Our general alertness is influenced by words we hear, scenes we see, and processes which require consciousness and interpretation of perceptions and which certainly are dependent on cortical activity...It appears very likely that corticoreticular projections are involved in these processes (p. 443).
>
> It is...common experience that increased attention and alertness are accompanied by an increased heart rate and often also other autonomic phenomena. This is easily explained by the general,

> ascending and descending, actions of the activating system (cf. axons with dichotomizing, ascending and descending branches in the reticular formation). (p. 442)."

The fact that the ascending and descending influences of the reticular activating system share common efferents suggests a method for measuring reticular involvement in cognitive processing; the measurement of task-evoked changes in the output of the descending activation system.

4.1 Pupillometric Measurement of Task-evoked Reticular Activation

Electrical measurement of reticular formation activity in man is a practical impossibility, as this structure consists of a dense network of cell bodies and fibers buried in the central core of the brainstem, extending from the medulla to the diencephalon. However, as indicated by Brodal, the state of reticular activity can be assessed indirectly from its effect on the autonomic reflexes, which are controlled by autonomic nuclei in adjoining regions of the brainstem. These nuclei have intimate connections with the reticular core. Of the autonomic functions regulated in this area, the pupillary system appears to be uniquely well-suited for the assessment of reticular activity during information processing in man.

Pupillary dilation has been used as a primary measure of reticular activation in neurophysiological research since the pioneering investigations of reticular function by Moruzzi and his co-workers (see Moruzzi, 1972, for a detailed summary). There are at least two reasons for this choice. First, the mechanical construction of the iris is such that small changes in either sympathetic or parasympathetic discharge may be discerned. Second, the central connections of the iridic musculature are intimately linked to nuclei controlled in large part by the reticular activating system.

Superimposed upon the tonic level of pupillary dilation, orderly phasic responses may be discerned during the performance of mental tasks. These task-evoked pupillary responses are time-locked to ongoing mental events. They are rapid in onset, being measurable at latencies of a few hundred milliseconds. They are relatively small in extent, reaching a maximum value of approximately .50 - .70 mm. at high processing loads, and terminate within a few hundred milliseconds of completion of cognitive processing. The amplitude of these task-evoked pupillary responses appears to be quantitatively related to the momentary level of processing demands imposed by a cognitive task.

In the fifteen years since pupillometric measurements were introduced to modern psychology (Hess and Polt, 1964), pupillary responses have been analysed in a variety of information-processing tasks. A number of cognitive domains have been investigated, including memory, perception, attention, language processing and reasoning. In each of these areas, quantitative relations between processing demands and the magnitude of the task evoked pupillary response have been established. The results of these experiments are summarized below.

4.2 Mental Arithmetic.

Mental arithmetic has been used as an example of a complex reasoning problem by several investigators interested in the pupillometric analysis of processing demands. Hess and Polt (1964), in their initial and influential article on pupillary signs of mental activity, measured pupillary diameter as 5 subjects solved 4 multiplication problems, ranging in difficulty from 7 X 8 to 16 X 23. For each of the subjects and each of the problems, pupillary diameter increased from the moment of problem presentation until the point of solution. Across subjects, the percentage dilation was perfectly ordered by presumed problem difficulty.

These results were subsequently replicated by Ahern (Ahern, 1978; Ahern & Beatty, 1979, in press). Three levels of problem difficulty were employed, ranging from multiplying pairs of 1-digit numbers to multiplying pairs of 2-digit numbers. In this task an initial dilation of approximately .15 mm accompanies the encoding and storage of the multiplicand. The second and major dilation follows presentation of the multiplier and continues through problem solution. Both the amplitude and latency of this latter dilation increase as a function of problem difficulty. In the most difficult condition, the response appears to asymptote at approximately .50 mm. An example of these task-evoked pupillary responses is shown in Figure 5.

Insert Figure 5 Here.

4.3 Short-term Memory.

The demands placed upon short-term memory formed an initial and enduring problem in the pupillometric study of information-processing load. Kahneman and Beatty (1966) presented the first pupillometric analysis of the processing demands encountered in a short-term memory task. Strings of 3 to 7 digits were auditorily presented at the rate of 1 per sec. After a two second pause, subjects were required to repeat the digit string at the same rate. Under these conditions, pupillary diameter increases in a linear fashion with the presentation of each digit, reaching a maximum in the pause preceding report. As digits are unloaded from memory during report, pupillary diameter decreases with each digit reported, reaching baseline levels after report of the final digit. The magnitude of the peak pupillary dilation during the pause between input and output is an increasing function of string length. In unpublished work, Kahneman and Beatty observed that if the subject were requested to repeat the string a second time immediately after reporting the final digit, the pupil immediately dilates to the peak diameter for that string and then decreases with each digit spoken until the entire string has been reported for the second time. Beatty and Kahneman (1966) demonstrated that a similar pupillary function is obtained when a string of items is recalled from long-term memory for report.

Pupillary dilation is also determined by the difficulty of the to-be-remembered information as measured by memory span for different types of items (Kahneman and Beatty, 1966). Three conditions were tested: recall of four digits, recall of four unrelated nouns, and transformation of a four digit string by adding one to each item. Steeper slopes were found for the more difficult item strings.

The capacity of short-term memory for strings of unrelated digits is approximately seven (Miller, 1956). Peavler (1974) measured the task-evoked pupillary response for strings of 5, 9 and 13 digits. During presentation of the strings pupillary diameter increases as an approximately linear function of memory load for digits 1 through 7. At the seventh or eighth digit, the pupillary response reached an asymptote; no further dilation was observed. Thus, the task-evoked pupillary response reflects increasing task demands only within the region of that function where adequate performance may be maintained.

4.4 Language Processing.

Several levels of language processing have been studied pupillometrically. At the most molecular level, Beatty and Wagoner (1978) used an experimental method developed by Posner and his colleagues (Posner & Mitchell, 1967; Posner & Boies, 1971) to study the visual encoding of single letters. In their first experiment subjects were required to judge whether or not a pair of visually-presented letters had the same name. Individual letters were presented in either upper or lower case type. Thus, two sorts of letter pairs could be judged to be the same by the name criterion. If both letters are presented in the same case (e.g. AA or aa) only the physical features of the letters need by analyzed to reach the correct judgment. If they differ in case (e.g. Aa or bB) then, in addition to analyzing the physical features, a second step of name code extraction must be performed. Although the task-evoked pupillary responses were small in this simple task (on the order of .1 mm), the pupil was sufficiently sensitive to indicate the extra processing step required for name code extraction. Significantly larger responses were obtained for letter pairs that differed in case.

In a second similar experiment, Beatty and Wagoner examined three levels of character encoding by requiring the use of a third, higher order category for classification (vowels and consonants). Some letter pairs could either be physically identical, identical in name, or identical in category membership (e.g. Ae or BK). Again, the task-evoked pupillary response differentiated the additional processing load required to perform the letter-matching task at each level.

Ahern (Ahern, 1978; Ahern & Beatty, 1979) undertook two experimental investigations involving language processing. The first of these experiments examined task-evoked pupillary responses in the perception and comprehension of words. Subjects were required to judge pairs of words as similar or different in meaning. The first word of each pair was drawn from either the easiest or the most difficult items of one of three psychometric vocabulary tests. The second word, presented two seconds later, was either a synonym of the first, or quite different in meaning. A dilation of approximately .1 mm followed the presentation of the easy target words, whereas the dilation to the difficult target words was twice as large. A second dilation followed the presentation of the comparison word, yielding pupillary dilations of .30 and .35 mm. respectively during the judgment period.

In the second experiment, Ahern (Ahern, 1978; Ahern & Beatty, 1979), employing Baddeley's Grammatical Reasoning Task (Baddeley, 1968), presented sentences of the form "A follows B" or "B precedes A" after which an exemplar "AB" or "BA" was given. The task was to determine whether the sentence correctly described the exemplar. Sentences differed in grammatical complexity, being active-positive, active-negative, passive-positive or passive-negative. Although these sentences differed in length, sentence duration was held constant by using computer presentation of digitized natural speech. In this task, increasing dilation was observed during the presentation of the sentence and the exemplar which peaked during the decision interval. The amplitude of these responses averaged approximately .40 mm and differed significantly as a function of grammatical complexity, with the longer, more complex sentences eliciting larger pupillary responses.

Wright and Kahneman (1971) have also applied pupillometric measurements in a sentence processing task. Subjects were presented with complex sentences of the form: "The qualified managing director was recently sensibly appointed by the expanding successful company." When required to repeat the sentence, the task-evoked pupillary response increased as the sentence was presented, and peaked during the retention interval (3 or 7 sec), reaching a maximum dilation of approximately .30 mm.

4.5 Perception and Attention.

Small but reliable pupillary dilations accompany the detection of both visual and acoustic signals at near threshold intensities. Hakerem and Sutton (1966) provided the first pupillometric analysis of processing load in perceptual detection. Subjects viewed a uniform visual field upon which brief increments in luminance could be imposed with the left eye as pupillary diameter of the right eye was measured. When the magnitude of the intensity increment was adjusted to yield 50% correct detection, all vestiges of the flash-induced light reflex disappeared. A clear pupillary dilation of approximately .10 mm was observed if and only if a presented target was detected. This dilation reflects the processing of detected targets. Although of small magnitude, these task-evoked responses to detection were highly reliable.

Beatty and Wagoner (1977) extended Hakerem and Sutton's finding to audition, using weak 100 msec 1kHz sinusoidal acoustic signals presented against a background of white noise with a probability of 0.50. After each trial, the subject rated his certainty that a target had or had not been presented (Green & Swets, 1966). For signal present trials, the magnitude of the task-evoked pupillary response was an increasing function of rated detection certainty, being largest for certain detections and smallest for certain rejections.

Task-evoked pupillary responses have also measured in perceptual discrimination tasks, in which a presented stimulus must be compared against memory and a judgment rendered. Kahneman and Beatty (1967) reported the first study of the pupillary response in pitch discrimination. The amplitude of the response to the comparison tone varied as an increasing function of discrimination difficulty, ranging from approximately .10 mm for easy to .20 mm for difficult discriminations.

4.6 Inter-task Comparisons.

The data summarized above consistently indicate that tasks which place large demands on the information processing system, judged behaviorally, subjectively or by an analysis of task requirements, elicit large task-evoked pupillary responses. Less demanding tasks elicit smaller responses. An intriguing possibility is that this pupillometric measurement of central nervous system activation associated with cognitive functions might provide a common metric for the assessment and comparison of information-processing load in tasks that differ substantially in their functional characteristics.

This possibility is strengthened by the finding that the magnitude of the task-evoked pupillary responses during cognitive processing is independent of baseline pupillary diameter over a physiologically reasonable range of values (Bradshaw, 1969, 1970; Kahneman and Beatty, 1967;

Kahneman, Beatty and Pollack, 1967). Therefore the absolute values of the task-evoked dilations reported from different laboratories for qualitatively different tasks may be meaningfully compared.

Insert Figure 6 Here.

Figure 6 presents such a quantitative comparison, giving the peak amplitude of the task-evoked pupillary response for each of the tasks detailed above, subject only to the constraint that the data are not confounded by the effects of overt responding.

The leftmost panel of Figure 6 presents peak dilations for short-term memory tasks. The data for short-term retention of digits are the average of the values obtained by Ahern (1978), Kahneman and Beatty (1966), Kahneman, Onuska and Wolman (1968), and Peavler (1974). The value for retention of 4 words is from Kahneman and Beatty (1966). The next panel summarizes the literature on language processing. The peak value for the letter matching task is the average of both experiments published by Beatty and Wagoner (1978). Sentence encoding-1 is from Wright and Kahneman (1971). Sentence encoding-2 is from a sentence repetition task recently completed by Beatty and Schluroff. All other values for language processing tasks are taken from Ahern (1978). Word encoding is the response to the presentation of the first word in the synonyms judgment task. The values for easy and hard word matching are the peak response during the judgment period following presentation of the second word in that task. The value for grammatical reasoning is the average of the four types of sentences in Baddeley's Grammatical Reasoning Task.

The third panel presents data from the mental multiplication task used as an example of complex reasoning. Only Ahern (1978) has presented task-evoked pupillary responses for this task which are necessary for comparative peak measurement. Multiplicand storage is the amplitude of the peak response to the first item in the multiplication task. The other three values are the peak amplitudes attained during problem solution.

The rightmost panel presents data for perceptual tasks. The visual detection data are from Hakerem and Sutton (1966) and the auditory detection data are from Beatty and Wagoner (1977). The discrimination data are taken from Kahneman and Beatty (1967).

Several points concerning these data deserve mention. First, the data are tolerant of the stringent demands placed upon them in comparing absolute dilation values across experiments. Usually rescaling of some sort is required for physiological data to remove individual differences in responsivity (Johnson and Lubin, 1972). No such rescaling was undertaken here. The data plotted are absolute peak dilations obtained from different groups of subjects performing a wide range of cognitive tasks under varying experimental conditions in different laboratories. Second, the data plotted in Figure 6 are internally consistent. No abnormally large or small values are present. Third, the ordering of these values corresponds closely to an ordering of these tasks using other criteria of information-processing load. The short-term memory tasks cover a large range of values, depending on the number of items held for recall. Similarly in language processing the sentence comprehension tasks yield large pupillary dilations whereas the simpler word and letter matching tasks elicit much smaller values. The mental multiplication tasks also span a wide range of values, each appropriate to the difficulty of the particular problem. Finally, the perceptual tasks, which behavioral techniques indicate impose negligible processing load, are associated with small task-evoked pupillary responses. Taken as a whole, these data indicate that the task-evoked pupillary response may serve as an independent, physiological indicator of the momentary level of information processing load imposed upon the nervous system by cognitive tasks.

5.0 Toward a Neurophysiological Theory of Processing Resources

Several facts have emerged from this review of the neurophysiology of cognitive processing. The first is that the cerebral cortex is not an undifferentiated structure, but rather that it is functionally specialized in anatomically definable regions. This conclusion is supported by both measures of regional cerebral glucose utilization and blood flow, each an indicator of the local rate of oxidative metabolism resulting from information transactions in cortical tissue. It is further corroborated by the neuropsychological studies of the effects of restricted brain lesions, although that literature was not reviewed here. These data give strong support to the idea that at the level of cerebral cortex multiple resource theory provides an appropriate description of the structure of the cognitive system. Furthermore, the neurophysiological data provides an initial definition for these multiple resources that may be confirmed by appropriate behavioral tests.

Second, these cortical resources have been shown to be expandable, in the sense suggested by Kahneman (1973). Cortical resources are not continuously active, but rather are allocated as appropriate given the cognitive tasks facing the person. Furthermore, there is a real physical cost associated with resource utilization; cortical resources require biologically large amounts of energy to process information. This finding is in accord with the analysis of resource utilization and processing costs offered by Navon and Gopher (1979) in their microeconomic model of cognitive processing. The fact that processing costs are real suggests that cortical resource utilization might be a carefully regulated biological process.

The third finding is that at the level of the brainstem activation systems, there is strong support for the idea of a general processing resource as Kahneman (1973) proposed. Across a wide range of qualitatively different information processing tasks, the amplitude of the task-evoked pupillary response appears to provide a consistant and sensitive indication of the aggregate demand for cortical processing resources.

Thus human neurophysiology has provided some answers to the major questions facing cognitive theory. But these data also suggest that the presently available models of the structure of processing resources may not be completely adequate. The problem is that there exists compelling neurophysiological evidence for both multiple processing resources at the level of cerebral cortex and a single general resource at the level of the brainstem. How are these two sorts of data to be reconciled? One answer to this question was suggested by an analogy proposed by Kahneman and Beatty in 1967:

> "The frequent use of the concept of processing load in the present paper has been guided by a simple analogy: consider a houseful of electrical devices that are variously put in operation by manual switches or by their own internal governor systems. The total amperage demanded by the entire system at any one time may easily be read on an appropriate electrical instrument outside the house. Processing load is here construed as analogous to an aggregate demand for power, and there is ground for the hope that the pupil may function as a useful approximation to the relevant measuring device. (Kahneman and Beatty, 1967, page 104)"

This analogy suggests a neurophysiological framework for a structural theory of processing resources. Specific, multiple processing resources of modern cognitive theory are identified with regionally restricted cortical areas of specialized function. Thus, in addition to behavioral data obtained in timesharing experiments (Navon and Gopher, 1979), neurobiological data may also be employed to identify and characterize probable specific processing resources. Neuropsychological data detailing cognitive deficits following restricted cortical lesions appear to be relevant (Walsh, 1978; Hecaen and

Albert, 1978; Heilman and Valenstein, 1979). Tasks which show impairment as the result of a specific, restricted brain lesion may be assumed to draw heavily upon at least one cortical resource in common. Similarly, recently developed methods of metabolic mapping including regional cerebral blood flow (Ingvar, 1976) and cerebral glucose utilization (Sokolov, 1979; Phelps et al., 1979) provide information concerning the regional involvement of specific cortical regions in neurologically normal individuals performing a variety of cognitive tasks. Thus, the analogy suggests the view that restricted, functionally similar regions of cerebral cortex may serve as specialized specific processing resources, a perspective that not only is compatible with modern multiple resource theory, but provides several sources of enriching, convergent information.

The analogy also suggests that the general processing resource may be identified with the aggregate demand from these regional cortical resources for activation from the reticular core. For reasons outlined above, the task-evoked pupillary response may be interpreted as an indicator of the momentary extent of cortically controlled reticular activation elicited during the execution of a particular cognitive process. It should be noted that the pupil is not the only indicator of reticular activation. Similar findings have been reported in both the autonomic (Kahneman, Tursky, Shapiro, and Crider, 1969) and skeletal-motor (Tuttle, 1924) systems. Within the skeletal-motor system, the amplitude of monosynaptic stretch reflexes increases dramatically during cognitive processing, an effect that has been attributed to an increase in excitability of the gamma-efferent system, Note 3), which is a classical indicator of reticular outflow

(Granit, 1970). This line of argument leads to the conclusion that cognitive processing elicits dynamic changes in the momentary level of activity in the ascending (and descending; see Brodal, 1981) reticular activating system in a manner quantitatively related to within-task and between-task variations in processing load. What is then required is a definition of the mechanisms linking the specialized forebrain processing resources and the general reticular activating system.

The mechanism by which the reticular core may activate or facilitate information processing in cerebral cortex has remained unclear for several decades following Moruzzi and Magoun's (1949) discovery of the behavioral activating properties of the reticular core. However, recent data from Schiebel (1980) and Skinner (Skinner, 1979) have suggested a solution to this problem. They propose that the reticular system functions to modulate an inhibitory gating mechanism controlling thalamocortical communication in the forebrain. The essence of this suggestion is that reticular activation momentarily expands the processing capacity of the forebrain by facilitating communication in the intrinsic thalamocortical pathways. Such a mechanism is in close accord with the electrical analogy for processing resources and processing load originally proposed by Kahneman and Beatty (1967) to account for the regularity of their pupillometric data.

This neurophysiological framework for a hierarchical theory of processing resources both provides a perspective for further theory development and suggests relevant sources of data from both the disciplines of cognitive psychology and the neurosciences. Perhaps most importantly, this framework for a theory, even without elaboration, generates a number of testable hypotheses. Pairs of tasks may be expected to interfere with each other as they compete for cortical resources. Therefore, either the neuropsychological literature or the developing regional brain metabolism literature may be used to predict pairs of tasks that are most likely to be mutually interfering. This prediction implies that these neurophysiological findings are useful in specifying structure-specific (Wickens, 1979) processing resources. Further, since the model is hierarchical, one might expect that timeshared tasks using markedly different cortical resources might nonetheless be limited by the total reticular activation available and therefore show interference. Such hypotheses should be testable by pupillometric methods.

The immediate importance of this approach is that it suggests ways in which the strong pupillometric evidence for a general concept of processing load is compatible with behavioral evidence for multiple, specialized processing resources. The resulting model is hierarchical in form and may provide a means of integrating the several, apparently discordant facts that distress simpler models.

In addition this perspective also may offer longer range advantages. The types of theories spawned from a neurobehavioral framework are likely to be richer in several ways than are theories developed from more limited perspectives. Such a theory is likely to be relevant to questions posed in both the psychological laboratory and the neurological clinic. By employing a wider range of material in theory construction, unprofitable theoretical branches are more likely to receive early pruning and interesting possibilities are more likely to be nurtured rather than discarded.

Finally, a neurobehavioral framework brings to the question of human information processing resources a richness of empirical data that is commensurate with the complexity of the theoretical problem.

References

Ahern, S.K. Activation and intelligence: Pupillometric correlates of individual differences in cognitive abilities. Doctoral dissertation, University of California at Los Angeles, 1978.

Ahern, S.K., and Beatty, J. Physiological signs of information processing vary with intelligence. Science, 1979, 205, 1289-1292.

Ahern, S.K., and Beatty, J. Physiological evidence that demand for processing capacity varies with intelligence. In M. Friedman, J.P. Das, and N. O'Connor (Eds.), Intelligence and learning. New York: Plenum, in press.

Bachelard, H.S. How does glucose enter brain cells? In D.H. Ingvar and N.A. Lassen (Eds.), Brain work: The coupling of function, metabolism and blood flow in the brain. Copenhagen: Munksgaard, 1975.

Baddeley, A.D. A three-minute reasoning test based on grammatical transformation. Psychonomic Science, 1968, 10, 341-342.

Beatty, J., and Kahneman, D. Pupillary changes in two memory tasks. Psychonomic Science, 1966, 5, 371-372.

Beatty, J., and Wagoner, B.L. Activation and signal detection: A pupillometric analysis. UCLA Human Neurophyiology Laboratory Technical

Report, 1977.

Beatty, J., and Wagoner, B.L. Pupillometric signs of brain activation vary with level of cognitive processing. Science, 1978, 199, 1216-1218.

Bradshaw, J.L. Pupil size and drug state in a reaction time task. Psychonomic Science, 1970, 18, 112-113.

Bradshaw, J.L. Background light intensity and the pupillary response in a reaction time task. Psychonomic Science, 1969, 14, 271-272.

Brodal, A. Neurological anatomy in relation to clinical medicine, 3rd Edition. New York: Oxford University Press, 1981.

Foit, A., Larsen, B., Hattori, S., Skinhoj, E., and Lassen, N.A. Cortical activation during somatosensory stimulation and voluntary movement in man: A regional cerebral blood flow study. Electroencephalography and Clinical Neurophysiology, 1980, 50, 426-436.

Granit, R. The basis of motor control. London: Academic Press, 1970.

Green, D.M., and Swets, J.A. Signal detection theory and psychophysics. New York: Wiley, 1966.

Hakerem, G., and Sutton, S. Pupillary response at visual threshold. Nature, 1966, 212, 485-486.

Hecaen, H., and Albert, M.L. Human neuropsychology. New York: Wiley, 1978.

Heilman, K.M., and Valenstein, E. Clinical neuropsychology. New York: Oxford, 1979.

Hess, E.H., and Polt, J.H. Pupil size in relation to mental activity during simple problem solving. Science, 1964, 143, 1190-1192.

Ingvar, D.H. Functional landscapes of the dominant hemisphere. Brain Research, 1976, 107, 181-197.

Ingvar, D.H., and Lassen, N. Regulation of cerebral blood flow. In H.E. Himwich (Ed.), Brain metabolism and cerebral disorders. New York: Spectrum, 1976.

Johnson, L.C., and Lubin, A. On planning psychophysiological experiments: Design, measurement and analysis. In N.S. Greenfield and R.A. Sternbach (Eds.), Handbook of psychophysiology. New York: Hold, Rinehard and Winston, 1972.

Kahneman, D. Attention and effort. Englewood Cliffs, N.J.: Prentice-Hall, 1973.

Kahneman, D., and Beatty, J. Pupil diameter and load on memory. Science, 1966, 154, 1583-1585.

Kahneman, D., and Beatty, J. Pupillary responses in a pitch-discrimination task. Perception and Psychophysics. 1967, 2, 101-105.

Kahneman, D., Beatty, J., and Pollack, I. Perceptual deficit during a mental task. Science, 1967, 157, 218-219.

Kahneman, D., Onuska, L., and Wolman, R.E. Effects of grouping ont he pupillary response in a short-term memory task. Quarterly Journal of Experimental Psychology, 1968, 20, 309-311.

Kahneman, D., Tursky, B., Shapiro, D., and Crider, A. Pupillary, heart rate and skin resistance changes during a mental task. Journal of Experimental Psychology, 1969, 79, 164-167.

Krnjevic, K. Coupling of neuronal metabolism and electrical activity. In D.H. Ingvar and N.A. Lassen (Eds.), Brain work: The coupling of function, metabolism and blood flow in the brain. Copenhagen: Munksgaard, 1975.

Larsen, B., Skinhoj, E., and Lassen, N.A. Variations in regional cortical blood flow in the right and left hemispheres during automatic speech. Brain, 1978, 101, 193-209.

Lashley, K.S. Brain mechanisms and intelligence: A quantitative study of injuries to the brain. Chicago: University of Chicago Press, 1929.

Lassen, N.A., Ingvar, D.H., and Skinhoj, E. Brain function and blood flow. Scientific American, 1978, 239, 62-71.

Lassen, N.A., Roland, P.E., Larsen, B., Melamed, E., and Soh, K. Mapping of human cerebral functions: A study of the regional cerebral blood flow pattern during rest, its reproducibility and the activations seen during basic sensory and motor functions. Acta Neurologica Scandinavia, Supplement 64, 1977, 262-263.

Luria, A.R. The working brain. New York: Basis Books, 1973.

Mazziotta, J.C., Phelps, M.E., Miller, J., and Kuhl, D.E. The metabolic map of the unstimulated resting human brain in transverse and coronal section. 32nd Annual Meeting of the American Academy of Neurology, New Orleans, 1980.

Melamed, E., and Larsen, B. Regional cerebral blood flow during voluntary conjugate eye movements in man. Acta Neurologica Scandinavia, Supplement 64, 1977, 530-531.

Miller, G.A. The magical number seven, plus or minus two: Some limits of our capacity for processing information. Psychological Review, 1956, 63, 81-97.

Moray, N. Where is capacity limited? A survey and a model. Acta Psychologica, 1967, 27, 84-92.

Moruzzi, G. The sleep-waking cycle. Reviews of Physiology: Biochemistry and Experimental Pharmacology. New York: Springer-Verlag, 1972.

Moruzzi, G., and Magoun, H.W. Brain stem reticular formation and activation of the EEG. Electroencephalography and Clinical Neurophysiology, 1949, 1, 455-473.

Navon, D., and Gopher, D. Task difficulty, resources and dual-task performance. In R. Nickerson (Ed.), Attention and performance VIII, Hillsdale, N.J.: Erlbaum, 1980.

Navon, D., and Gopher, D. On the economy of the human processing system. Psychological Review, 1979, 86, 214-255.

Norman, D.A., and Bobrow, D.G. On data-limited and resource-limited processes. Cognitive Psychology, 1975, 7, 44-64.

Orgagozo, J.M., and Larsen, B. Activation of the supplementary motor area during voluntary movement in man suggests it works as a supramotor area. Science, 1979, 206, 847-850.

Peavler, W.S. Individual differences in pupil size and performance. In M. Janisse (Ed.), Pupillary dynamics and behavior. New York: Plenum, 1974.

Phelps, M.E., Kuhl, D.E., and Mazziotta, J.C. Metabolic mapping of the brain's response to visual stimulation: Studies in Humans. Science, 1981, 211, 1445-1448.

Phelps, M.E., Mazziotta, J.C., Kuhl, D.E., Nuwer, M., Packwood, J., and Metter, J. Tomographic mapping of the metabolic response of the visual cortex to stimulation and deprivation: Studies in man. 32th Annual Meeting of the American Academy of Neurology, New Orleans, 1980.

Phelps, M.E., Huang, S.C., Hoffman, E.J., Selin, C., Sokoloff, L., and Kuhl, D.E. Tomographic measurement of local cerebral glucose metabolic rate in humans with (F-18)2-Fluoro-2-Deoxy-D-Glucose: Validation of method. Annals of Neurology, 1979, 6, 371-388.

Posner, M.I., and Boies, S.J. Components of attention. Psychological Review, 1971, 76, 391-405.

Posner, M.I., and Mitchell, R.F. Chronometric analysis of classification. Psychological Review, 1967, 74, 392-409.

Roland, P.E., Skinhoj, E., Larsen, B., and Lassen, N.A. The role of different cortical areas in the organization of voluntary movements in man: A regional cerebral blood flow study. Acta Neurologica Scandinavia, Supplement 64, 1977, 542-543.

Sanders, A.F. Some remarks on mental load. In N. Moray (Ed.), Mental workload: Theory and measurement. New York: Plenum, 1979.

Scheibel, A.B. Anatomical and physiological substrates of arousal. In J.A. Hobson and M.A.B. Brazier (Eds.), The reticular formation revisited: Specifying function for a nonspecific system. New York: Raven, 1980.

Skinner, J.E. A neurophysiological model for regulation of sensory input to cerebral cortex. In D.A. Otto (Ed.), Multidisciplinary perspectives in event-related brain potential research. Washington, D.C.: U.S. Environmental Protection Agency, 1979.

Sokolov, L. Mapping of local cerebral functional activity by measurement of local cerebral glucose utilization with (14-C)deoxyglucose. Brain, 1979, 102, 653-668.

Tuttle, W.W. The effect of attention or mental activity on the patellar tendon reflex. Journal of Experimental Psychology, 1924, 7, 401-419.

Walsh, K.W. Neuropsychology: A clinical approach. Edinburgh: Churchill Livingstone, 1978.

Wickens, C.D. The structure of attentional resources. In R.S. Nickerson (Ed.), Attention and performance VIII, Hillsdale, N.J.: Erlbaum, 1980.

Wickens, C.D. Human operator workload. In N. Moray (Ed.), Mental workload: Theory and measurement. New York: Plenum, 1979.

Wright, P., and Kahneman, D. Evidence for alternative strategies of sentence retention. Quarterly Journal of Experimental Psychology, 1971, 23, 197-213.

Figure Captions

Figure 1. Two dimensional FDG scans of a human subject 40 minutes after intraveneous injection. High concentrations of FDG may be seen in the brain and heart. The three images are taken three angles with respect to frontal view (60 degrees right, frontal, and 60 degrees left). (From Mazziotta et al., 1980)

Figure 2. Cerebral metabolic rate for glucose in visual cortex reflects visual information processing demands. (A) Sketches from actual brain slices at approximately the level of the cross-sectional tomographic images. The blackened areas at the posterior of the brain indicate approximate locations of primary (PVC) and association (AVC) visual cortices. (B) Metabolic activity in visual cortex is low when the eyes are closed and patched. (B) Stimulation by white light. Increasing CMRGlc is indicate by increasing shades of gray, with black being the highest. Higher metabolic rates may be seen for PVC and AVC (arrows) with stimulation. (D) Metabolic response of the visual cortex while viewing a complex scene. (From Phelps, Kuhl and Mazziotta, 1981.)

Figure 3. Mean increase in the glucose metabolic rate of the PVC and AVC in homonymous hemianopsia and with increasing complexity of visual stimulation. Data are percentage increases over eyes-closed controls. The rate of increase for AVC is relatively greater with complex patterned stimuli. (From Phelps et al., 1980.)

Figure 4. A schematic view of the left cerebral hemisphere with major functional areas indicated.

Figure 5. Averaged task-evoked pupillary responses obtained in mental multiplication. Responses for three levels of problem difficulty are plotted. The magnitude of the response during the solution period increases as a function of processing demands. (From Ahern, 1978.)

Figure 6. Peak amplitudes of the task-evoked pupillary responses obtained in a rage of qualitatively different cognitive tasks, arranged by type of task. The pupillary response provides a reasonable ordering of tasks on the basis of presumed processing load. See text for further details.

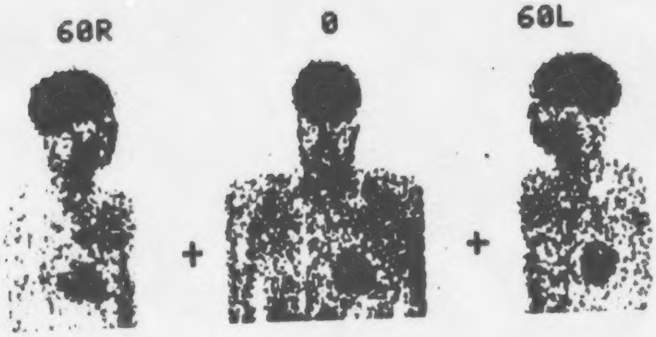

169

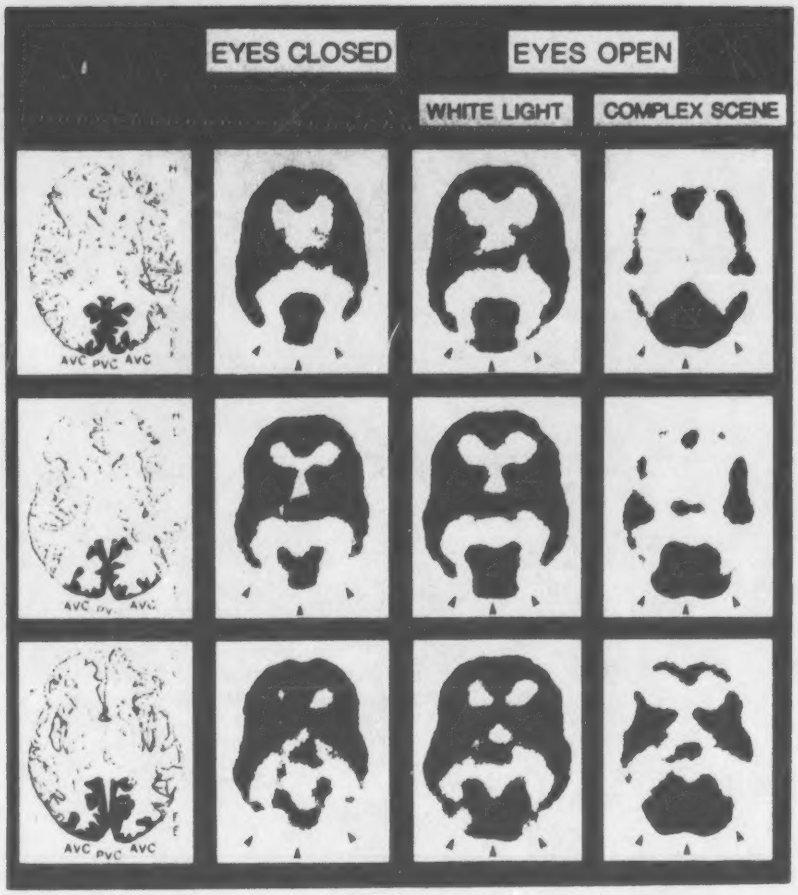
EYES CLOSED
EYES OPEN
WHITE LIGHT
COMPLEX SCENE
AVC PVC AVC
AVC PVC AVC
AVC PVC AVC

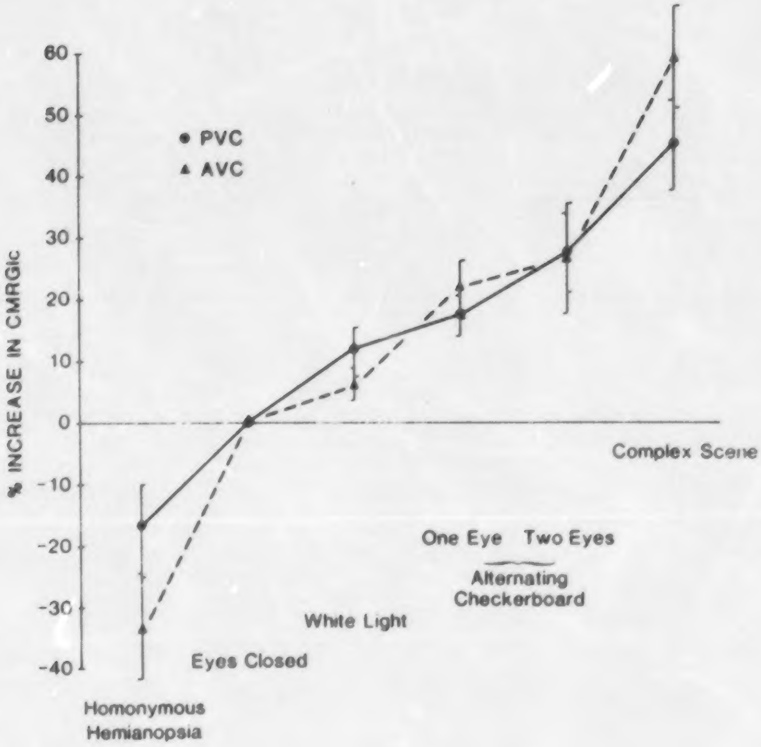

% INCREASE IN CMRGlc
60
50
40
30
20
10
0
-10
-20
-30
-40
● PVC
▲ AVC
Homonymous Hemianopsia
Eyes Closed
White Light
One Eye
Two Eyes
Alternating Checkerboard
Complex Scene

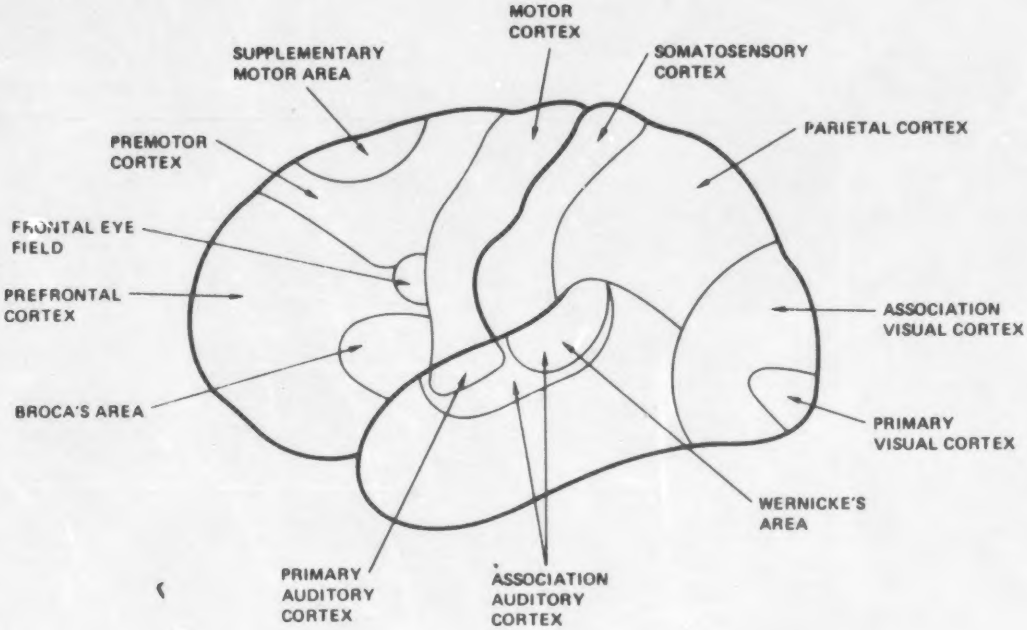
MOTOR CORTEX
SUPPLEMENTARY MOTOR AREA
SOMATOSENSORY CORTEX
PREMOTOR CORTEX
PARIETAL CORTEX
FRONTAL EYE FIELD
PREFRONTAL CORTEX
ASSOCIATION VISUAL CORTEX
BROCA'S AREA
PRIMARY VISUAL CORTEX
WERNICKE'S AREA
PRIMARY AUDITORY CORTEX
ASSOCIATION AUDITORY CORTEX

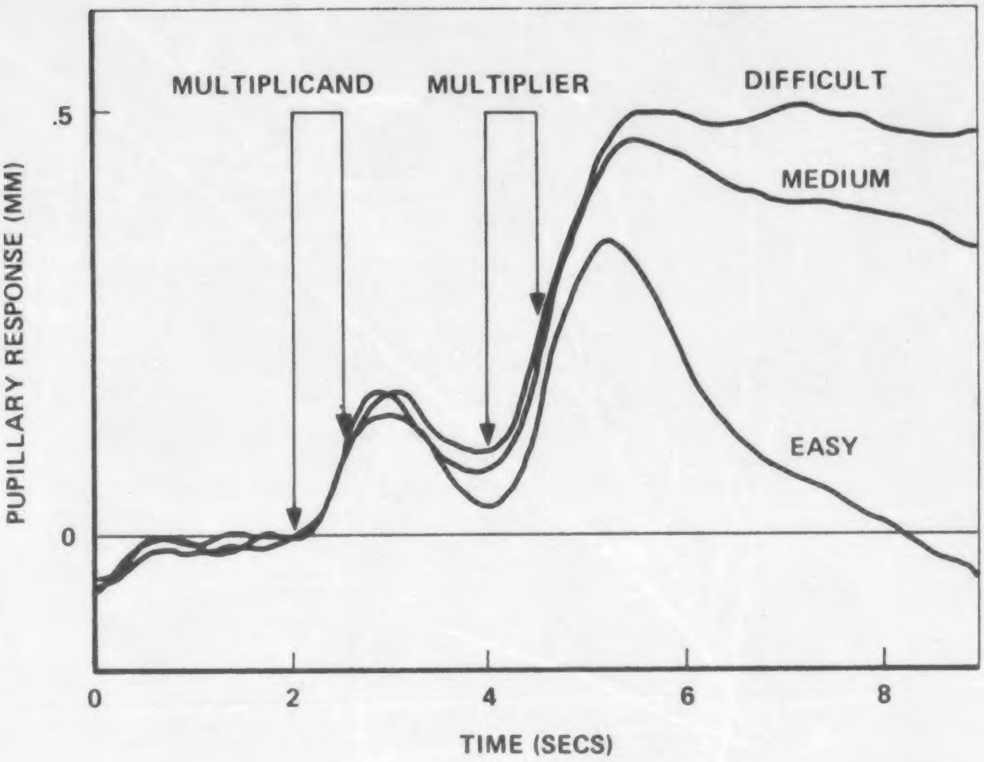
MULTIPLICAND
MULTIPLIER
DIFFICULT
MEDIUM
EASY
.5
0
PUPILLARY RESPONSE (MM)
0
2
4
6
8
TIME (SECS)

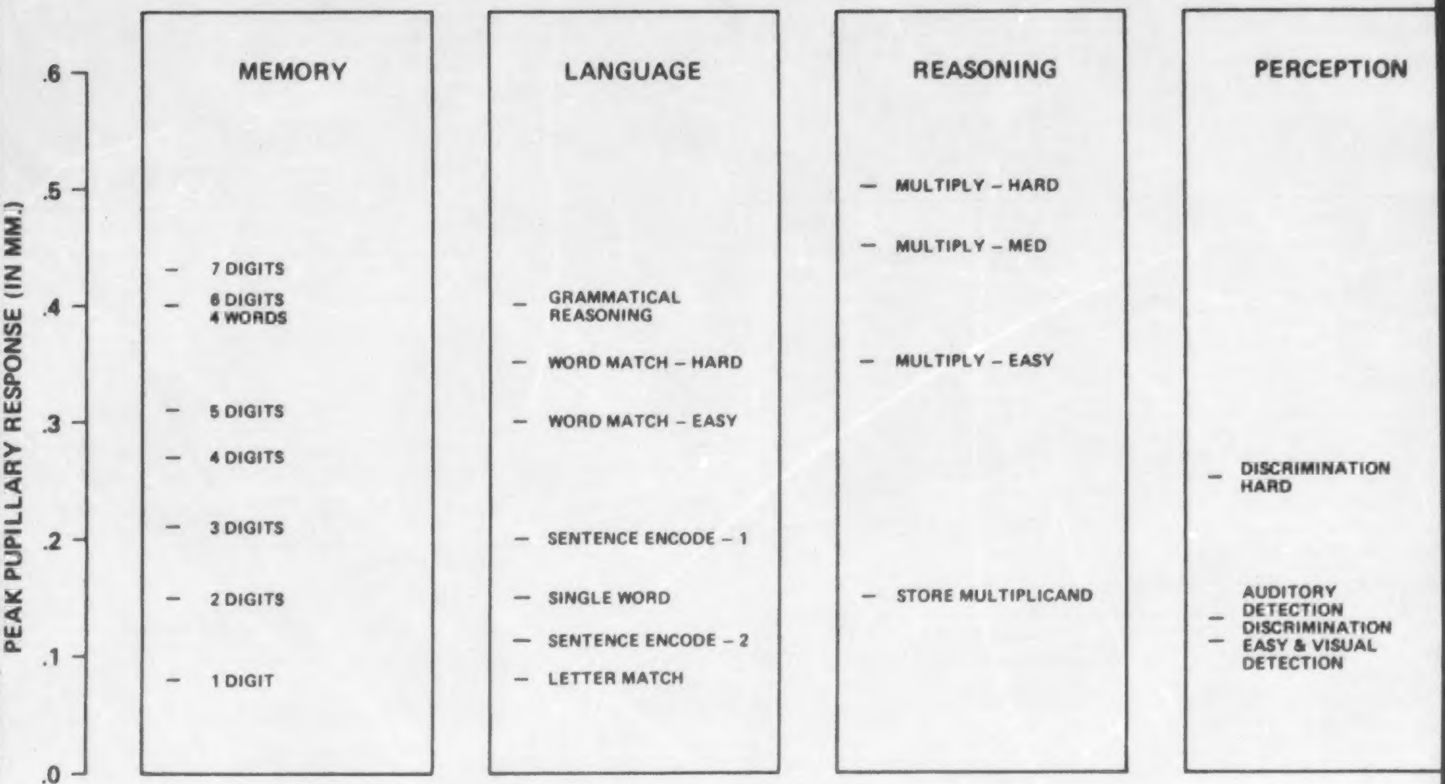
PEAK PUPILLARY RESPONSE (IN MM.)
.6
.5
.4
.3
.2
.1
.0
MEMORY
7 DIGITS
6 DIGITS
4 WORDS
5 DIGITS
4 DIGITS
3 DIGITS
2 DIGITS
1 DIGIT
LANGUAGE
GRAMMATICAL REASONING
WORD MATCH – HARD
WORD MATCH – EASY
SENTENCE ENCODE – 1
SINGLE WORD
SENTENCE ENCODE – 2
LETTER MATCH
REASONING
MULTIPLY – HARD
MULTIPLY – MED
MULTIPLY – EASY
STORE MULTIPLICAND
PERCEPTION
DISCRIMINATION HARD
AUDITORY DETECTION
DISCRIMINATION EASY & VISUAL DETECTION

ELECTROPHYSIOLOGICAL STUDIES OF ATTENTION IN MAN

STEVEN A. HILLYARD

I would like to describe some electrophysiological studies of attentional mechanisms in man. Our goal has been to gain insight into the structure of attentional processes by recording the time-locked electrical fields that arise from the brain during the processing of sensory information.

What most of us in physiological research would ultimately like to achieve is to specify exactly what brain circuits and what brain structures are involved in the processing of information and in selecting among stimulus inputs. For this type of understanding, however, we'll probably have to wait a few decades, or even longer because we're still very far from being able to localize specific information processing activities in the human brain. Scalp electrodes, in particular, provide a very limited perspective on the underlying patterns of neural activity and the brain structures that are active. However, the recording of neural correlates from the scalp as markers of underlying attentional and cognitive processes, does, I think, provide information that will help us understand more about the nature of those processes.

I think that the recording of neural correlates from the scalp can be of use in several ways. Our general strategy has been to record evoked potentials or event related potentials (ERPs), as they're called, from the scalp in carefully structured behavioral circumstances where we have an idea of what perceptual or attentional processes are active. Having identified a wave as a marker of a more or less well-defined attentional process, we can go on to new tasks and get some idea of the timing of that process in a different processing sequence, and see how it interacts with other processes that have also been marked with ERPs. Specifically, by looking at the latency of ERPs, we can get information on the question of which stages or what levels of processing participate in attentional selections - the old issue of early selection versus late selection. We can also get information about how much processing is accorded to unattended or irrelevant material to which the subject is not making behavioral responses. Typically, a stimulus situation involves many more stimuli that one can respond to behaviorally, and the behavioral readout provides a very limited view of how a person is processing a multi-dimensional stimulus array. By examining the ERPs which are triggered by all the stimuli that are being presented, one can gain a more complete picture of how each type of stimulus is being processed. So, the ERP technique provides another way of looking at perceptual processes in the brain which has advantages and disadvantages in relation to purely behavioral approaches; as a long-term goal we're trying to bring these two lines of research into correspondence with one another.

- - - - - - - -

INSERT SLIDE 1

- - - - - - - -

The first slide shows the general approach, with electrodes placed on the head of a typical subject and the raw, amplified EEG shown above. When a warning flash and a signal click are presented you don't see very much in the raw EEG, because these stimuli evoked ERPs in the brain which are just a few microvolts in amplitude while the on-going EEG "noise" is of the order of tens of microvolts.

So, in order to see these tiny evoked brain signals, we have to use a signal averaging computer and summate the ERPs to large numbers of signals, typically between tens and hundreds of signals. The end result is the averaged event-related potential, and if we just look at the ERP to the click signal, we can see a series of potential oscillations, some sixteen consistent waves, each with its own label. These waves are quite reproducible from one person to the next, and a number of them have been associated with anatomically localized neural activity. Within the first ten milliseconds after a click is a series of oscillations which can be localized to specific brainstem auditory pathway nuclei as the message ascends through the brainstem on its way to the cortex (Stockard & Rossiter, 1977; Starr, 1978). The wave V at a latency of 6 msec comes from the mid-brain, probably the inferior colliculus, and wave VI probably from the vicinity of the medial geniculate body. Somewhere in the latency range of 10-15 msec the message arrives at the cortex and gives rise to a whole series of late cortical waves. It is in the late, presumably cortical activity beyond 50 msec where the ERPs associated with selective attention are seen.

Early experiments in our laboratory (Picton & Hillyard, 1974) were designed to examine the old descending inhibition theory of auditory attention (Hernandez-Peon, 1966) which proposed that attention to non-auditory stimuli resulted in a gating of auditory transmission in the brainstem via efferent inhibitory pathways. Contrary to this hypothesis, we found that the amplitude of the brainstem evoked potential to clicks did not vary under conditions of attention versus those of inattention. Thus, we obtained no evidence for gating of auditory input at the brainstem level.

The most consistent changes in ERPs with selective auditory attention were observed in a broad negative wave which begins at about 50 sec and peaks at 100 msec; this wave is termed N_1, since it was the first negative wave that was discovered by Davis and associates. Most of my talk will be concerned with how this negative ERP varies as a function of what a person is attending to. The second slide shows the standard paradigm that we've been using for a number of year to elicit these negative ERPs related to attention. (Hillyard et al., 1973, 1978).

INSERT SLIDE 2

In this task stimuli such as tone pips are delivered to different sensory channels; in this case we present high frequency tones to one ear and low frequency tones to the other ear. These tones are given in random order to the two ears at a rapid rate, typically at intervals of 200 to 400 milliseconds or about three beeps per second. This places a very high load of information on the subject, making it very difficult to attend to all the stimuli that are occurring. The subject is asked to focus attention on the tones in one ear and to detect occasional targets in that ear which are difficult to discriminate from the non-targets. In this type of experiment, for instance, 10% of the tones (the targets) might be of slightly longer duration and sound a little louder than the non-targets. As the person listens to the tones in one ear and tries to detect those targets, the selective attention effect is seen in the ERP beginning at about 60 or 70 milliseconds after the tone onset. These are scalp recordings from the vertex, in the midline of the head. If you look at the ERPs to left ear tones when the left ear is being attended, those tones evoke large negative waves; when attention is shifted over the right ear, the left ear tones evoke small negative waves. The reverse happens to the ERPs elicited by the right ear tones. In general, the N_1 amplitude is high to attended-channel tones and is reduced to unattended tones. In conditions where neither ear is attended (reading a book), the N_1 has an intermediate amplitude.

We've come to think of this attention-related ERP in terms of the difference between the attended and the unattended channels; this difference (between the solid and dotted lines) is actually a broad negativity that usually begins before the N_1 peak and extends considerably after that peak. We used to think of this ERP as a simple enhancement of the N_1 wave, an enlargement of the evoked potential that is always elicited by a tone. More recently, however, we've come to think of it as a new, "endogeneous" negative wave that's added on the N_1 of the evoked potential under conditions of selective attention (Hansen & Hillyard, 1980). This ERP has recently been called the "processing negativity" by some workers in the field (Näätänen & Michie, 1979), but I usually call it simply the "N_1 effect."

Now, I'd like to give a brief historical resume of some of the factors which have been shown to modulate this attention effect (reviewed in Hillyard et al., 1978; Hillyard & Picton, 1979). In particular, the stimuli have to be delivered at a fairly high rate; if the tones are delivered to the two ears at a rate of one per second or less, there is no difference between the ERPs to attended and unattended tones. It appears that there has to be a sufficient load of information to engage the subject's selective attention. We also attempted to find out what kinds of sensory channels could be used to produce this negative ERP effect. In the experiment just described, the sensory channels were defined as the two ears, with the stimuli also differing in pitch. We later found that similar ERP changes are seen if the channels are defined simply as two spatial locations (tones of the same frequency from two spatial locations) or if the channels are defined as two different frequencies presented from the same location in space. The N_1 effect is also present if the stimuli are of the same pitch and at the same location but differ in intensity. That is, whenever two channels of stimuli differ in any simple physical attribute that can be easily and rapidly analyzed, and attention is focussed on one of the channels, the attended stimuli will elicit this enhanced processing negativity. Since this ERP effect had certain parallels with Broadbent's (1970, 1971) concept of a "stimulus set" or "filtering" type of attention, we suggested that the processing negativity was correlated with the additional processing of sensory input after it had been selected by a Broadbent style filter or stimulus set mechanism. Another property of this ERP effect which is consistent with Broadbent's concept is that the negativity is elicited by any stimulus that belongs to the attended channel, whether it's relevant to the task or not. Even though the subject has to respond to only some stimuli in a channel, all the stimuli belonging to that channel elicit the enhanced negativity.

We've also done experiments in which we looked at the amplitude of processing negativity as a function of how subjects distribute their attention among two or more auditory channels (Hink *et al.*, 1977). For instance, when subjects are asked to divide their attention between the two ears and listen for targets in both ears instead of one, the ERP amplitude is intermediate between the attended N_1 level and the unattended N_1 level seen during focused attention. Thus, when attention is focused on one channel, a large amplitude N_1 is elicited by all of the attended stimuli while the ignored channel stimuli elicit a small N_1; when attention is divided between the two ears, however, the amplitude is almost exactly intermediate between these two values. This implies that the total output of negativity over a given time period is going to remain rather constant. This constancy of the total negativity across conditions of focused and divided attention has been replicated by several laboratories (Parasuraman, 1978; Okita, 1979) and it appears that this negativity may represent a type of capacity-limited processing system. While this notion remains speculative, it is clear from the ERP data that the processing negativity reflects the allocation of attention to different stimulus channels and perhaps (as described below) to different stimulus attributes as well.

A recent experiment done in our laboratory by David Woods and colleagues examined the properties of the selective "filters" which allow us to focus attention on one spoken message in a noisy environment (e.g., in the classic cocktail party) while ignoring competing voices. In this simplified version of the cocktail party experiment, subjects wore earphones and listened to either a male voice reading a novel in one ear or a female voice reading another novel in the other ear. The study was designed to assess the specificity of the attentional filtering process using ERPs. We were interested in determining the nature of the information which was admitted to the attended ear and/or rejected from the inattended ear. The behavioral data on this phenomenon have been the focus of a classic debate. Originally, an all-or-none filter was proposed which presumably blocked out the entire message in the unattended ear. Lawson (1966), however, challenged this concept by demonstrating that reaction times were equivalent to tones presented to the shadowed and unshadowed ears. Indeed, it was shown that certain types of information did get through from the unattended channel (e.g., frequency of the speaker's voice, after Treisman; reviewed in Broadbent, 1971).

In this experiment, we recorded ERPs to four different kinds of "probe" stimuli that were superimposed on the speech messages in each ear, to see how each would be "filtered" by the attentional process. These probes were superimposed on the speech in either ear at a rate of about one per second. The probes were: (1) a tone pip at the fundamental frequency of the speaker's voice (about 100 Hz for a male voice and 200 Hz for the female voice); (2) a tone pip at the second formant frequency of the speakers voice (about 1000 and 1100 for the male and female voices, respectively); (3) the spoken word "ah"; and (4) the spoken word "but", the latter two in the actual voices of the speakers in each ear. These speech probes were produced by computer which made an A-D conversion of the person saying "ah" and "but" and superimposed these sound on the continuous speech.

The subject's task in one condition was to listen to the spoken message in one ear at a time and try to remember its content for a later questionnaire. The probe stimuli themselves were irrelevant. It was a standard speech monitoring situation--listen to this voice and ignore that voice. In a second condition, however, we required the subject to shadow, that is, to repeat the message in one ear aloud. Since the results were very similar for the shadowing and monitoring conditions, the ERP data in Slide 3 are collapsed over both conditions. We wanted to look at ERPs during shadowing in particular, because there's a report in the literature (Robinson & Sabat, 1975) which claims that when people shadow speech on one ear, they completely block out information in the other ear; that is, stimuli in the rejected ear elicited no cortical ERPs. We wanted to see if we could replicate that effect, which could have been caused artifactually by the masking of the unattended sounds by the shadower' voice. To avoid this problem of masking by the subject's voice, we presented the speech messages against a loud white noise background and had the subject whisper the message.

INSERT SLIDE 3

The effects of attention on the ERPs to the four different kinds of probe stimuli are shown in Slide 3. There's very little attention effect for the fundamental frequency tone; that is, very little difference between the ERPs to probes on the attended and unattended channels. There is also little ERP differentiation in the N_1 latency zone for the second formant probe; this is a high-pitched tone pip which stands out perceptually in both attended and rejected voices and also elicits the largest ERPs in both channels. This is what one would expect from the behavioral data in dichotic listening situations. In contrast, the ERPs to the speech sounds that are superimposed on the voices show strong effects of attentional selection. Both the "ah" sound and the "but" sound elicit a broad processing negativity when they occur in the attended ear.

The slide shows only ERPs at a central midline electrode, but lateral scalp recordings revealed that the second phase of the processing negativity is asymmetrically distributed on the scalp, being larger over the left cerebral hemisphere. Such asymmetries are not seen during attention to tone pips or non-verbal material (Hillyard & Woods, 1979). This suggests that the processing negativity which overlaps the N_1 wave reflects post-selection processing that is distributed in the brain according to the nature of the stimulus material to be analyzed; in the case of speech, of course, the left hemisphere should be the more actively engaged.

The main point I want to make with these data, however, is that the "filter" that selects one speech message in preference to another has rather specific properties; it's not simply admitting all the information within a particular frequency zone or in a particular ear, but rather it's selecting according to more complex properties of the speech sounds. The human speech selection system is a fairly sophisticated filter, if you will, tuned to the higher-order, patterned properties of the spoken message. Only stimuli having the appropriate sensory configuration are admitted for further processing.

The specificity of attentional filtering in the visual system has been investigated by Russ Harter and associates of the University of North Carolina. (Harter & Previc, 1978). When attention is focused upon a particular locus along a sensory dimension, not only is information preferentially admitted at that locus but there is also a zone of loci around the attended point which receive preferential processing. That is, if we consider a simple "searchlight" model of attention, the searchlight's beam will have a certain diameter or width. Harter's group has tried to define this "bandwidth" of the attentional channel that selects for a particular size (spatial frequency) of a checkerboard. Nine different sizes of checkerboards were presented in random order, and the subject was instructed to

183

attend to one particular size and press a button every time it occurred. The main comparison was between the ERP to the attended size checkerboard minus the ERP to that stimulus when it was not being attended. Again, the attention effect was a broad negativity, in this case distributed over occipital scalp regions, in contrast with auditory attention where the N_1 wave is more anteriorly distributed. Looking at difference waves across all check sizes, Harter & Previc found that the largest negativity was elicited by the attended size, with a progressive decline in amplitude for more distant check sizes. This function might be considered like a "roll-off" curve of the attentional selection process, where the ERP for each stimulus is plotted as a function of distance from the attended sensory locus. Harter's work illustrates another way in which ERPs have been exploited to define the detailed quantitative properties of an attentional process.

The overall picture of how visual attention affects ERPs, however, is far from clear. Paying attention to different sensory dimensions like color, position in space, orientation, or more complex patterns is associated with a different ERP configuration in each case.

This contrasts with auditory attention where (as long as the attended and unattended sounds belong to separate channels) the attended stimuli elicit a broad processing negativity regardless of the nature of the sounds. Thus, it is conceivable that the visual system is associated with a wider variety of attentional mechanisms than the auditory system.

Question: If you were to switch around the relevant stimuli rapidly; in other words, just give one or two trials, then give the subject new instructions or a new signal, so that the first stimulus is now irrelevant, would you get a different spectrum of negativity?

Yes. I think you would get a different result altogether. This effect is related to what Mike Posner was talking about--that sustained attention to a repeating source involves different attentional process than does switching attention on every trial. We've made that comparison in some pilot studies on attending to different visual locations; we found that when a subject sustains attention to a repetitive flashing light at one location, those flashes elicit an enlarged N_1 (N150) component (e.g., Van Voorhis & Hillyard, 1977). In contrast, when we did a variant of the Posner paradigm (Posner et al., 1978) where the subject switches attention on every trial, we didn't get the enhanced negative ERP at all. This ERP difference between sustained vs. rapidly switchable attention suggests that we are dealing with different varieties of attention, but I don't know yet how to characterize these differences.

Question: I was thinking that if you can get set for a long block of trials, you can sort of set up a gate and the selection becomes more automatic. In the other case, you've got to stop and do a lot more processing for each stimulus.

I think that's a good interpretation. The data are consistent with that idea. There are much larger ERP differences for sustained attention. However, I'm not sure whether these attention effects, which begin at 80-100 msec in the visual system are signs of an early gating process. In the auditory system, the ERP effects have properties which are very close to what one would expect from a Broadbent/Treisman type of filter. That is, when attention is directed towards stimuli having a simple physical attribute, all stimuli which share that attribute show equivalently large ERPs. The "processing negativity" seems to have the same dynamic properties as the Broadbent-style early selection system. But, how early is early, I don't know

Question: Does the intensity of the stimuli affect the N_1 wave?

Yes. Intensity has a strong effect. In general, the size of the N_1 wave, the processing negativity, is increased when stimulus intensities are reduced. This contrasts with the effect of lowering intensity on the evoked or exogenous components, which become smaller for fainter stimuli.

I would now like to describe some very recent experiments by Jon Hansen in our laboratory, which look into the question of how attention is allocated to the different attributes of multi-dimensional auditory stimuli. This experiment addresses an issue that Ann Treisman has been working on lately (Treisman _et_ _al_., 1977)--namely, when people attend to a stimulus that has two attributes like color and shape, or pitch and location, are those two attributes processed in parallel, independently, or is each stimulus analyzed from the beginning as a particular conjunction of those attributes. In other words, when you're looking for red letter N, do you process that stimulus immediately as a conjunction of attributes, a red N, or do you process it for its redness and "Nness" attributes separately at an initial stage, later conjoining the two attributes in a subsequent stage (e.g., of focal attention). These are two alternative models for what happens when you're attending to stimuli that have two attributes, and Hansen's experiments translate this question into ERP terms, making predictions about the waveforms that one might expect under each model.

The experimental paradigm is as follows: there are two sensory dimensions, pitch and location, with two levels of each dimension for a total of four stimuli; tone pips of two different pitches occurred at each of two locations. The subject is wearing earphones, and high and low pitched stimuli (1500 and 600 Hz) are given in random order to both left and right ears. The subject's task on each run is to listen to one of these four dual-attribute tones and to try and detect occasional occurrences of targets which are of a longer duration. This is a very difficult task, requiring the subject to discriminate pitch and location attributes as well as tone duration and to press a button when all these attributes are appropriate to the targets.

The main experimental question was the following: when you attend to one of these four combinations of pitch and location, do you process that tone as a conjunction of attributes, as a whole "object", or do you process each of the two attributes separately? Slide 4 shows some idealized waveforms of how the processing negativity might look under each of these alternatives. If attention is focused upon a conjunction of pitch-location attributes from the beginning, that stimulus should elicit an enlarged processing negativity, whereas each of the other three stimuli should elicit a much lower processing negativity. On the other hand, if the two attributes are processed independently, the amount of negativity elicited by each tone should vary according to how many attributes it shares with the attended pitch-location combination (Slide 4A, upper tracings). A stimulus which fulfilled both of the cue requirements, having the attended pitch as well as the attended location, should elicit the largest processing negativity. The tones that have one of attended attributes but not the other should have an intermediate amplitude, and tones with neither attribute would have the smallest negativity. For the hypothesis of independent attribute processing, the critical comparison is whether the ERP difference between tones having the attended attribute A versus the unattended attribute A varied as a function of whether attribute B was fulfilled or not. For instance, for tones of the attended pitch we can form a difference ERP between the attended and unattended locations (upper shaded area). If pitch and location are being processed independently, we would expect this difference ERP to the attended ear minus the unattended ear tones to be the same for tones of the unattended frequency as well (lower shaded area). That is, the differential processing of one attribute should not depend upon the presence of the other attribute under the hypothesis of independent, parallel processing.

- - - - - - -

INSERT SLIDE 4

- - - - - - -

One could also consider the possibility of a hybrid of these two models, wherein the attributes are first processed independently and then, when the attributes are combined, the tones come to be processed as conjunctions; thereafter, only the attended conjunction stimulus would receive further processing, and the other stimuli that failed to meet any one of the attended cue criteria would not receive further processing.

The next slide (5) shows the ERPs to the four stimulus alternatives (top tracings); these are ERPs to the shorter-duration, non-targets only. The top tracing is the ERP to stimuli that have the attended pitch and are in the attended location (ear). You see that the greatest processing negativity follows these tones, and there is quite a bit less processing negativity following each of the three other types of tones. That is, if a stimulus fails to fulfill either the pitch or the locational criterion, it does not elicit much processing negativity. This pattern of data is most consistent with conjoint processing model, although the fit is not perfect. These results are wholly inconsistent with the independent parallel processing model, however; this is particularly evident in the fact that waveforms 3 and 4 are nearly identical to one another in processing negativity. Those ERPs are both elicited by tones of the unattended pitch, but one of them occurs at the attended location while the other occurs in the opposite ear. This suggests that when a stimulus is recognized as having the "wrong" (unattended) pitch, there is no further selection on the basis of its location attribute. Now, if location and pitch were processed with total independence, there should be some ERP differentiation on the basis of location even if the stimulus failed to have the attended pitch attribute. This can also be illustrated in the "difference ERPs" shows below; difference wave 3 minus 4 shows the ERP differentiation between tones in the attended ear minus the unattended ear when the pitch cue is wrong. This is a very small difference. If we look at the difference wave for the location cue when the pitch cue is fulfilled, however (that's 1 minus 2) we see a large differential of processing negativity. If the two dimensions were processed in parallel, difference wave 1-2 should be the same as difference wave 3-4. That is, the processing of location should not depend on whether the pitch cue was fulfilled or not.

- - - - - - - -

INSERT SLIDE 5

- - - - - - - -

It appears, however, that the results are not entirely consistent with the conjoint processing model, because the ERP to the stimulus that has the correct (attended) pitch, but the wrong location is a bit larger than the ERP to tones which do not have the attended pitch attribute. That is, when a stimulus occurs at the wrong location, it is still processed to some extent if it has the right pitch. This suggests that pitch may have a kind of primacy over location as an attribute for this type of selection task. This effect is seen in the difference wave 2-4, showing some extra processing negativity following tones of the right pitch at the wrong location.

Question: How much training do the people have on this task?

Very little, these are mostly college students recruited for their first evoked potential experiment.

Question: Is it possible that these dual-attribute discriminations could become automated after extensive training at the task and that the N_1 component or processing negativity would change as a function of behavioral automization. You could train subjects to discriminate the four alternatives very rapidly, with very

189

fast reaction times, and see if the N_1 components change over that period of time.

We haven't looked into the question of learning and automization. I think that'd be an interesting problem to investigate. When behavior becomes automatized, do the ERPs reflect that change?

Question: I think that that's one of the question you've really got to answer, because if I understand Treisman's conjecture, it is not whether people initially process conjunctions, but whether or not they can come to process conjunctions after they've had sufficient experience with them.

Well, Treisman's two-stage model is not dependent upon special learning processes, as I understand it. But I don't know to what extent these effects are modifiable by training.

Further evidence on how people select multi-dimensional stimuli comes from other conditions in Hansen's experiment where the discriminability of one of the attributes was varied. The experiment was identical to the one just described, except that either the pitch or the location dimension was made less discriminable. To make the location discrimination more difficult, the interaural loudness balance was adjusted so that the tones were localized close together in space instead of in the separate ears. The ERPs to each of these four tones distinguished by "easy" pitch cues (600 vs. 1500 Hz) and "hard" location cues (localized close together) are shown at the right of Slide 6. The results suggest that paying attention to one of the four tones now involves two distinct phases of selection. At about 70 milliseconds after the stimulus there's an ERP differentiation along the "easy" dimension; that is, tones that have the appropriate (attended) pitch attribute start showing more negativity than do stimuli that have the wrong pitch. The solid and dashed lines are the ERPs to the tones having the attended pitch. If you're attending to the high pitch tones located slightly to the left of the midline say, and the

190

high pitch is easily discriminable from the low pitch, the high pitched tones start eliciting more negativity immediately. The ERP differentiation along the hard (location) dimension occurs subsequently; the processing negativity to the attended location becomes clearly larger only after a latency well beyond 100 msec. After about 400 msec, all three of the irrelevant tones produce equivalently small processing negativity.

INSERT SLIDE 6

This particular pattern of ERPs, I think, is consistent with a model of a hierarchical selection process (idealized in Fig. 4B). Why hierarchical? If you look at the lower ERPs, you see that once a stimulus fails to meet the easy criterion (i.e., has the unattended pitch), there is no further ERP differentiation along the hard (location) attribute. Once a stimulus fails to meet the pitch criterion for a target there's no differential processing according to what location it's in. A two-stage selection model would fit this data, with an initial selection based on the easy attribute and a later selection that's based on the more difficult attribute. Each stage is reflected in the prolongation or termination of the processing negativity. Put simply, if a stimulus fails by the easy pitch criterion, it's not processed at all for location.

Of course, the experiment also included another condition where pitch was the hard dimension (900 vs. 960 Hz) and location was the easy dimension, and in that case, there was a similar ERP pattern (shown in lower left waveforms of Slide 6). However, this latter pattern of ERPs suggests, as in the "easy-easy" condition, some minimal amount of processing of pitch information even when a tone fails to meet the location criterion. This data is again consistent with the idea that tones of the attended pitch in the wrong location (ear) receive a little more processing than tones of the unattended pitch.

In conclusion, what I've tried to illustrate here is that ERPs can provide useful evidence for the analysis of attentional processes, both to measure the timing of different stages of processing and, more generally, to classify different kinds of selection mechanisms.

I haven't talked about the P300 component yet, but if we were to look at the ERPs to the target stimuli that the subjects responded to, there would be a late positive wave at about 300-400 msec latency. The target tone ERPs would show an initial differentiation of processing negativity along the easy dimension, a second selection of negativity for the hard dimension, and then a final selection process by which the person decides it is a target; this latter decision would be associated with the positive P300. This ERP pattern suggests that the terminal selection of the target stimulus involves a very different kind of attentional mechanisms, associated with the P300, from the initial selections based on the other physical attributes of the stimulus. I think that the study of ERP configurations like these, or indeed, psychophysiological configurations in general, can help to differentiate among different types of attentional and cognitive processes.

Question: Does the P300 reflect a highly cognitive type of processing or does it have a strong perceptual weighting?

I'm not sure how to define the boundary between perceptual and cognitive processes. You can elicit P300s under very simple circumstances. For instance, if a repetitive sequence of stimuli is presented, any deviant stimulus that you have to react to differentially (e.g., make a counting response or a motor response) will trigger a P300 wave. I think that type of task might involve a very simple discriminative system without a high level of cognitive processing. Of course, the P300 only occurs subsequent to a sensory discrimination process and it does depend on memory storage of the relevant alternative stimuli and alternative plans of action. Such processes might be considered "cognitive". It is clear, incidently, that the P300 is not associated directly with the final motor output but with intermediate stages of stimulus evaluation and classification (see Donchin, 1979).

Question: I'm surprised at how long the processing negativity lasts, given that the tones are so short in duration. Why isn't there a rapid cut-off of negativity when duration is judged to be short and thus the stimulus "fails" on this third dimension?

Good question. The durations of the frequent, non-target tones are 50 milliseconds, while the targets are 100 milliseconds long. So you're asking why does the processing negativity extend out for several hundred milliseconds beyond the tone offsets. I would speculate that the duration discrimination (which is fairly difficult) is not made immediately, on-line, but occurs as the stimulus is being held in a sensory storage (or "echoic memory", as it's been called). That is, when a stimulus meets the proper pitch and location criteria, it is held in sensory memory for a further, more sustained analysis of its other properties, to see whether it's a longer duration target or not. The reaction times in this task are quite long, between 500 and 600 milliseconds. I would speculate that this sustained ERP reflects further processing of a stimulus "image" that is being held in memory for a longer period than its actual physical duration.

193

REFERENCES

Broadbent, D. E. Stimulus set and response set: Two kinds of selective attention. In: D. I. Mostofsky (Ed.) Attention: Contemporary Theory and Analysis. New York: Appleton-Century-Crofts, 1970.

Broadbent, D. E. Decision and Stress. New York: Academic Press, 1971.

Donchin, E. Event-related brain potentials: A tool in the study of human information processing. In: H. Begleiter (Ed.) Evoked Brain Potentials and Behavior Volume 2. New York: Plenum Press, 1979, 13-88.

Hansen, J. C. & Hillyard, S. A. Endogenous brain potentials associated with selective auditory attention. Electroencephalography and Clinical Neurophysiology, 1980, 49, 277-290.

Harter, M. R. & Previc, F. H. Size-specific information channels and selective attention: Visual evoked potentials and behavioral measures. Electroencephalography and Clinical Neurophysiology, 1978, 45, 628-640.

Hernandez-Peon, R. Physiological mechanisms in attention. In: R. W. Russell (Ed.) Frontiers in Physiological Psychology. New York: Academic Press, 1966.

Hillyard, S. A. & Picton, T.W. Event-related brain potentials and selective information processing in man. In: J. Desmedt (Ed.) Progress in Clinical Neurophysiology Volume 6 Cognitive Components in Cerebral Event-Related Potentials and Selective Attention. Basel: Karger, 1979, 1-50.

Hillyard, S. A. & Woods, D. L. Electrophysiological analysis of human brain function. In: M. S. Gazzaniga (Ed.) Handbook of Behavioral Neurobiology Volume 2 Neuropsychology. New York: Plenum, 1979, 345-378.

Hillyard, S. A., Hink, R. F., Schwent, V.L. & Picton, T. W. Electrical signs of selective attention in the human brain. Science, 1973, 182, 177-180.

Hillyard, S. A., Picton, T.W. & Regan, D. M. Sensation, perception and attention: Analysis using ERPs. In: E. Callaway, P. Tueting and S. Koslow (Eds.) Event-Related Brain Potentials in Man. New York: Academic Press, 1978, 223-322.

Hink, R. F., Van Voorhis, S. T., Hillyard, S.A. & Smith, T. S. The division of attention and the human auditory evoked potential. Neuropsychologia, 1977, 15, 597-605.

Lawson, E. A. Decisions concerning the rejected channel. Quarterly Journal of Experimental Psychology, 1966, 18, 260-265.

Näätänen, R. & Michie, P. T. Early selective attention effects on the evoked potential: A critical review and reinterpretation. Biological Psychology, 1979, 8, 81-136.

Okita, T. Event-related potentials and selective attention to auditory stimuli varying in pitch localization. Biological Psychology, 1979, 9(4), 271-284.

194

Parasuraman, R. Auditory evoked potentials and divided attention. *Psychophysiology*, 1978, *15*, 460-465.

Picton, T. W. &. Hillyard, S. A. Human auditory evoked potentials. II. Effects of attention. *Electroencephalography and Clinical Neurophysiology*, 1974, *36*, 191-200.

Posner, M. I., Nissen, M. J., & Ogden, W. C. Attended and unattended processing modes: The role of set for spatial location. In: H. L. Pick and E. Saltzman (Eds.) *Modes of Perceiving and Processing Information*. Hillsdale: Erlbaum, 1978.

Robinson, D. N. & Sabat, S. R. Elimination of auditory evoked responses during auditory shadowing. *Physiological Psychology*, 1975, *3*, 26-28.

Stockard, J. & Rossiter, V. Clinical and pathologic correlates of brainstem auditory response abnormalities. *Neurology*, 1977, *27*, 316-325.

Starr, A. Sensory evoked potentials in clinical disorders of the visual system. *Annual Review of Neuroscience*, 1978, *1*, 103-127.

Treisman, A. M., Sykes, M. & Gelade, G. Selective attention and stimulus integration. In: S. Dornic (Ed.) *Attention and Performance Volume VI*. New Jersey: Erlbaum, 1977, 333-363.

Van Voorhis, S. T. & Hillyard, S. A. Visual evoked potentials and selective attention to points in space. *Perception & Psychophysiology*, 1977, *22*, 54-62.

ONGOING EEG

AVERAGED EVENT-RELATED POTENTIALS

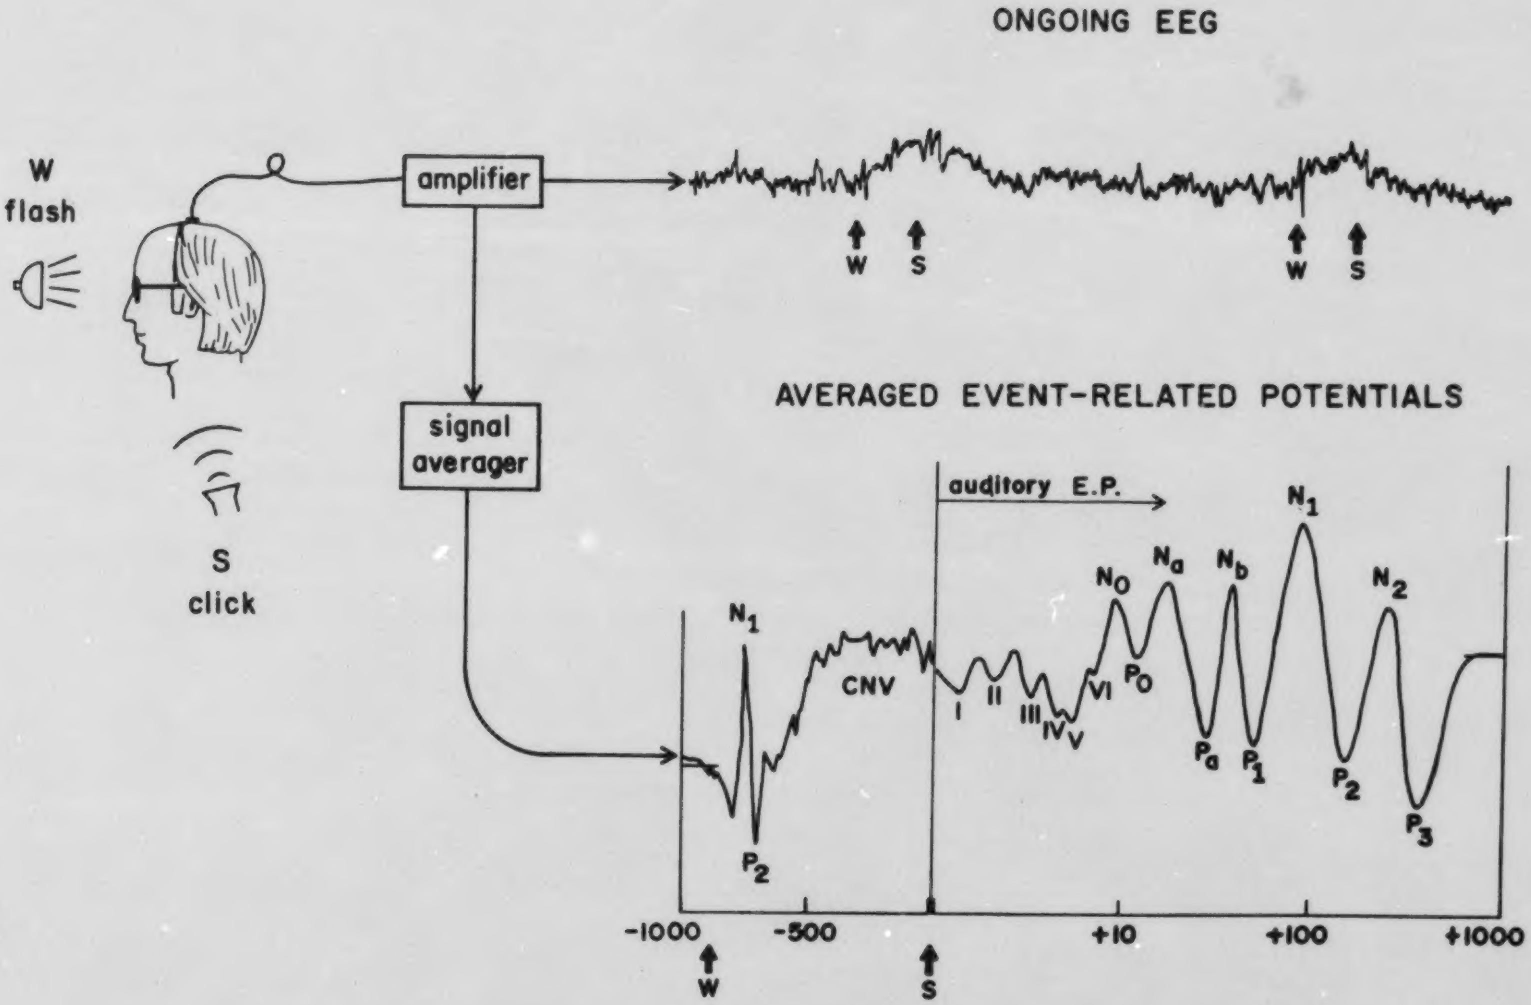
W
flash
S
click
amplifier
signal averager
W
S
W
S
auditory E.P.
N1
P2
CNV
I
II
III
IV
V
VI
N0
P0
Na
Pa
Nb
P1
N1
N2
P2
P3
-1000
-500
+10
+100
+1000
W
S

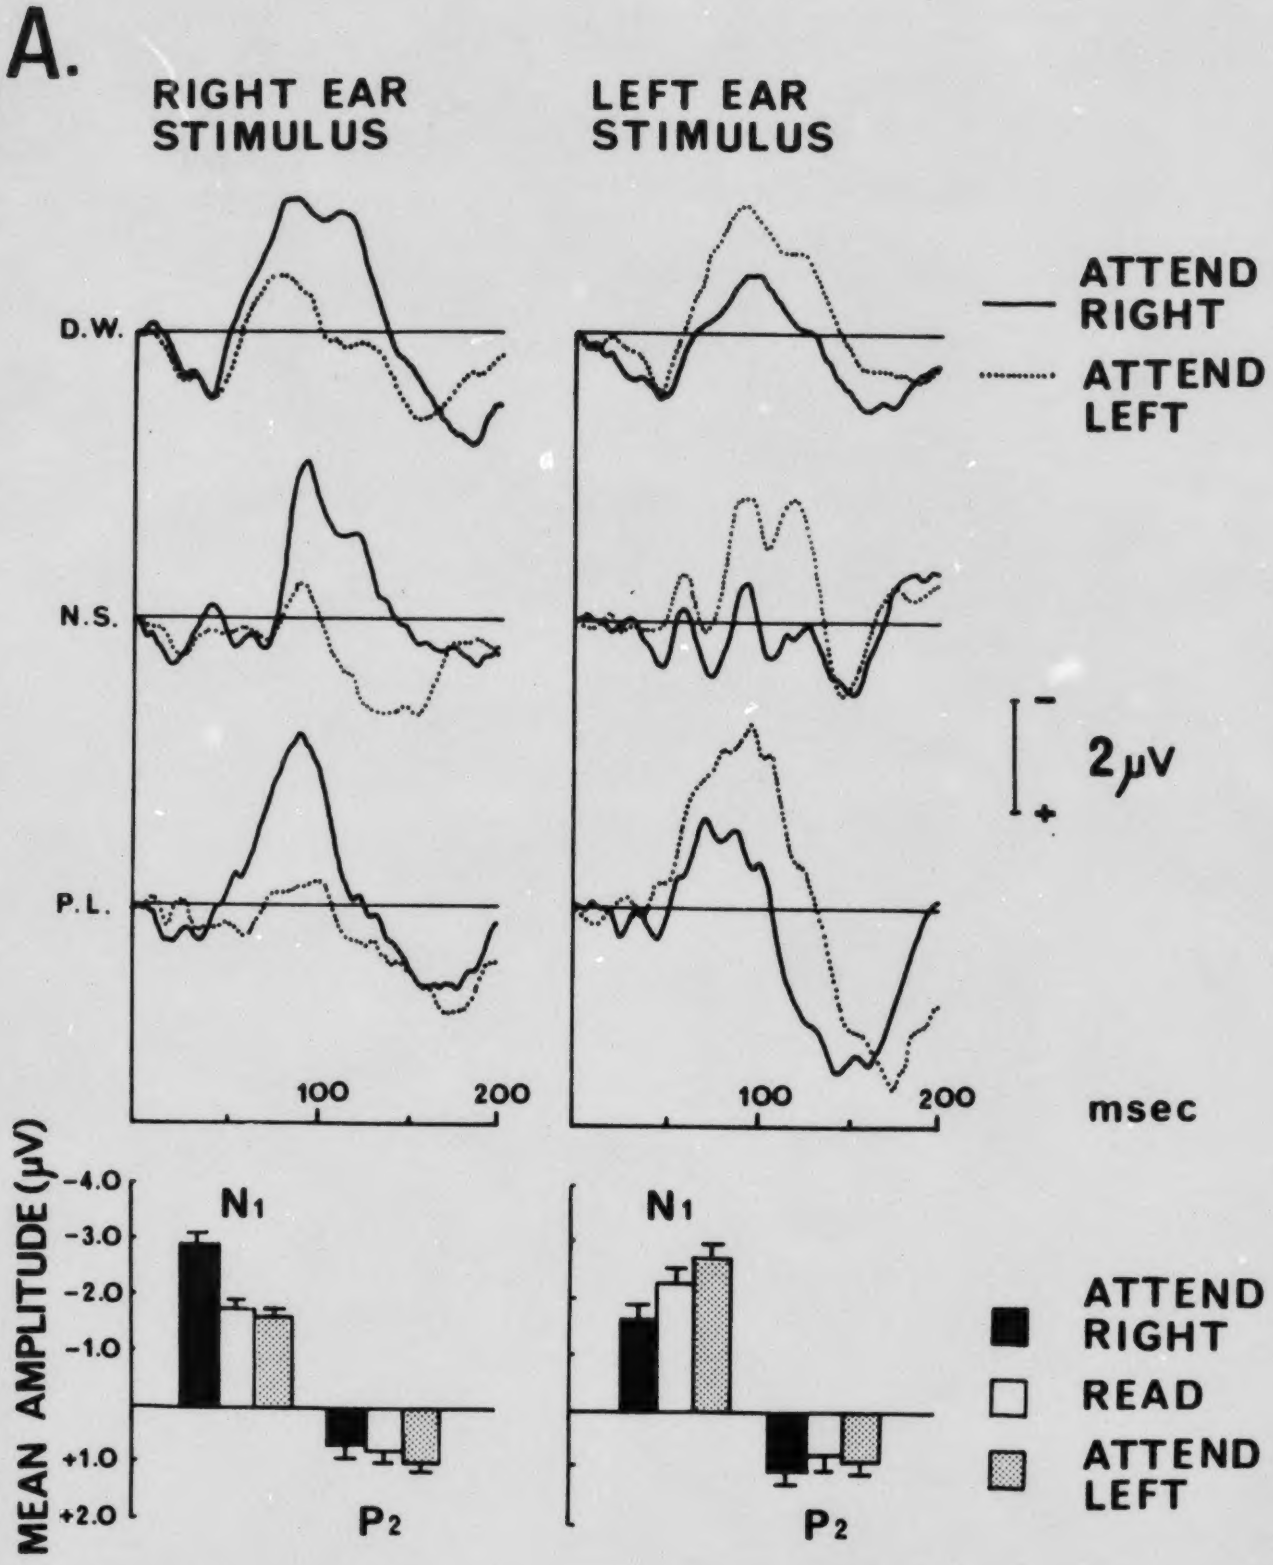
A.
RIGHT EAR STIMULUS
LEFT EAR STIMULUS
D.W.
N.S.
P.L.
ATTEND RIGHT
ATTEND LEFT
–
2µV
+
100
200
100
200
msec
MEAN AMPLITUDE (µV)
-4.0
-3.0
-2.0
-1.0
+1.0
+2.0
N1
P2
N1
P2
ATTEND RIGHT
READ
ATTEND LEFT

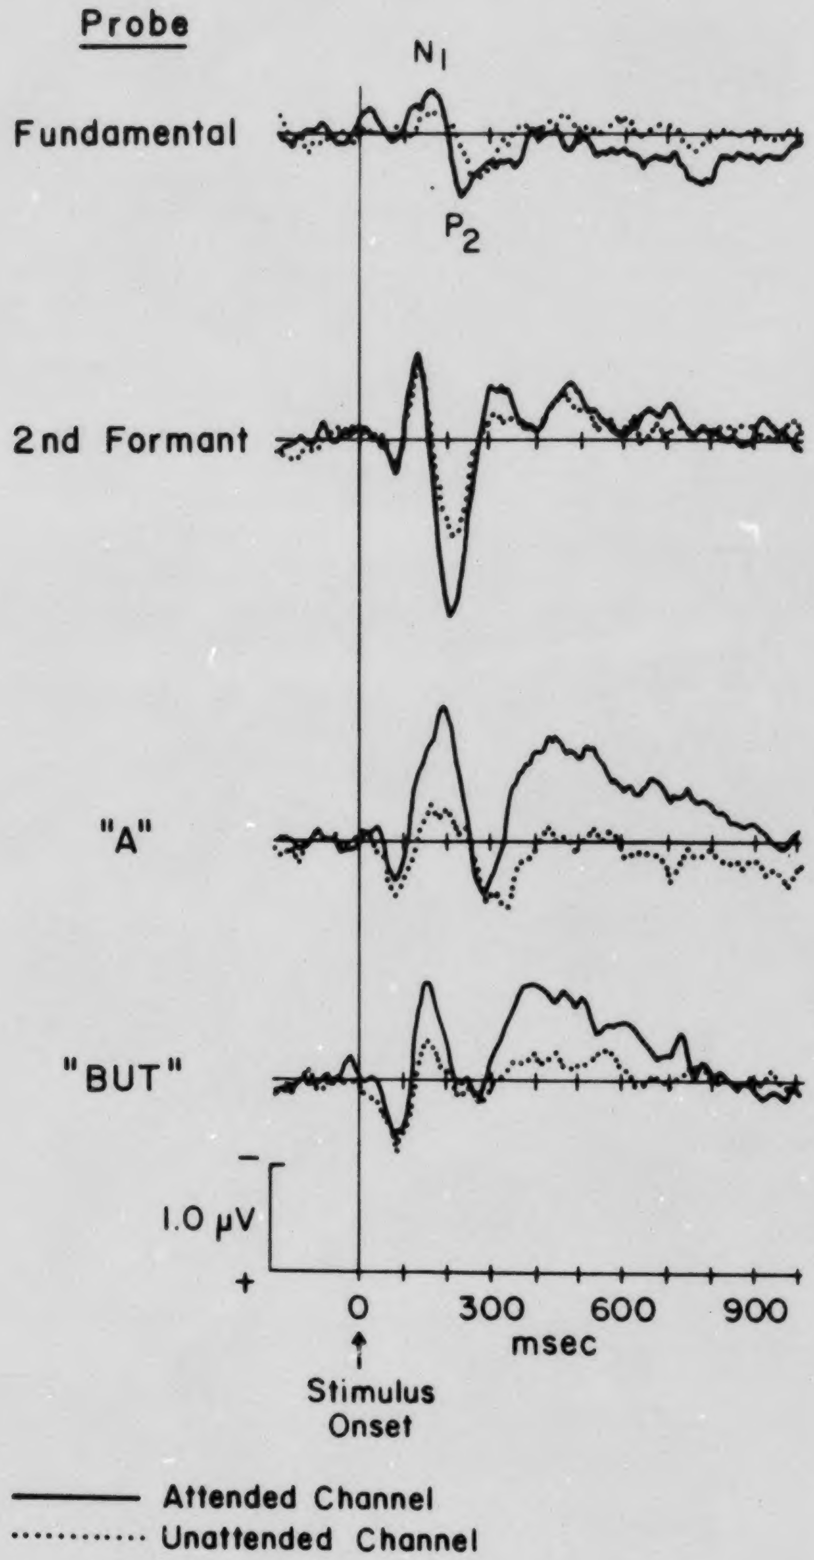
Probe
N1
Fundamental
P2
2nd Formant
"A"
"BUT"
−
1.0 µV
+
0
300
600
900
msec
Stimulus Onset
Attended Channel
Unattended Channel

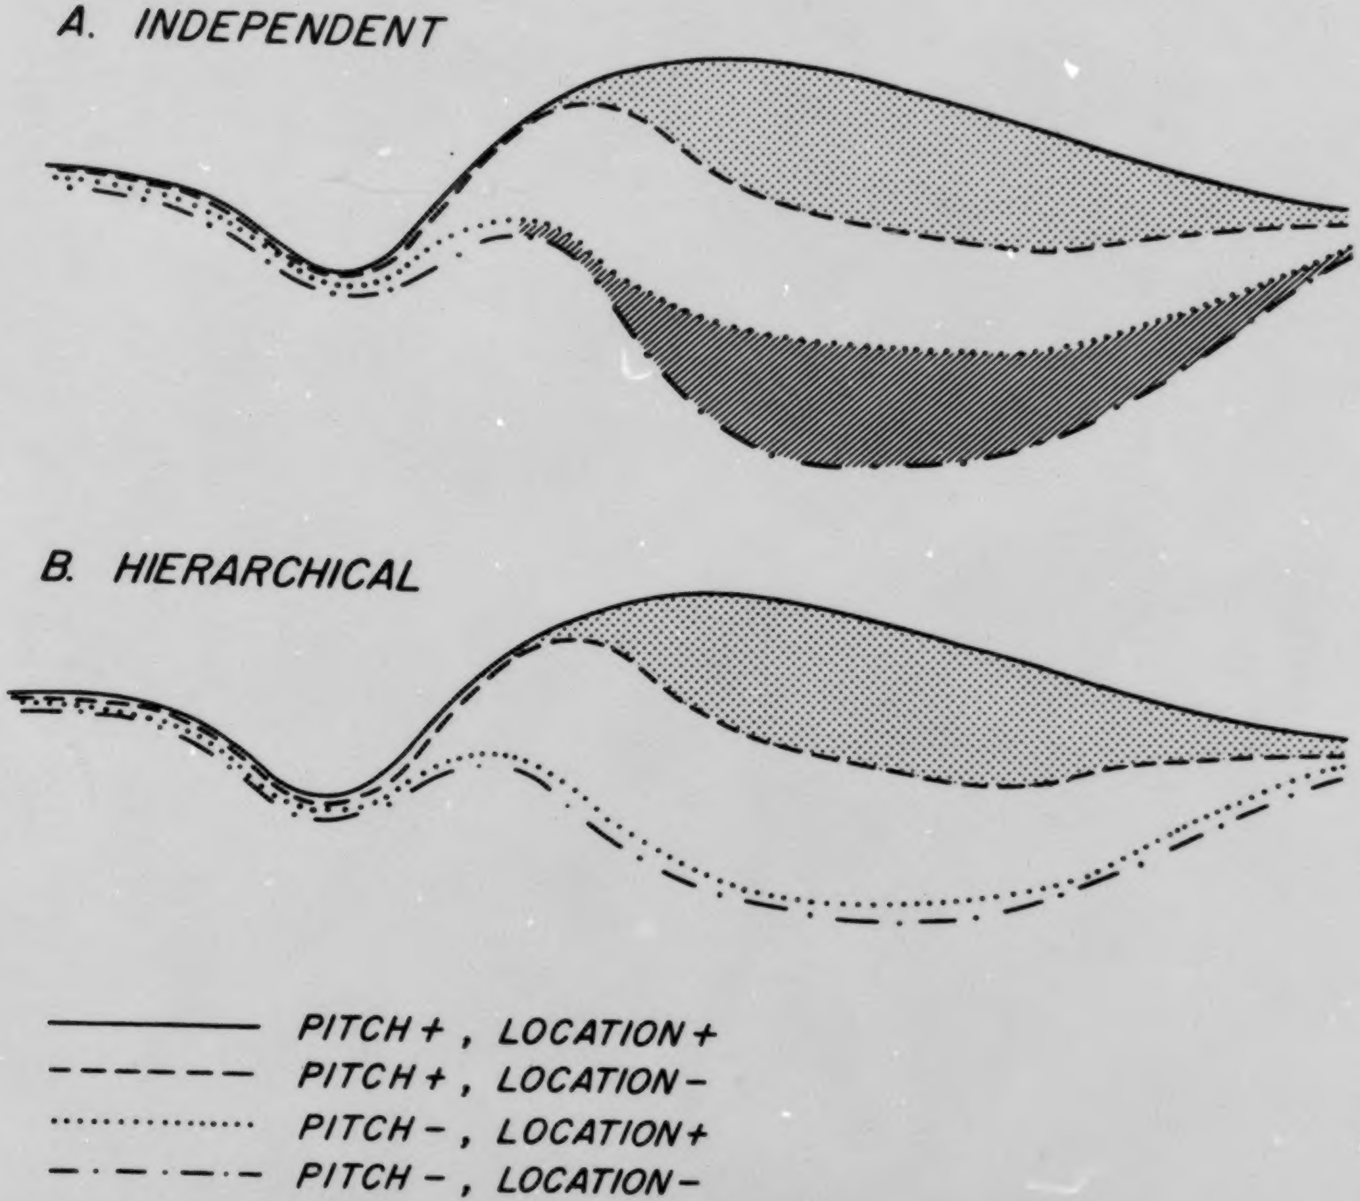
A. INDEPENDENT
B. HIERARCHICAL
PITCH + , LOCATION +
PITCH + , LOCATION −
PITCH − , LOCATION +
PITCH − , LOCATION −

PITCH EASY, LOCATION EASY
RAW WAVEFORMS (F_z)

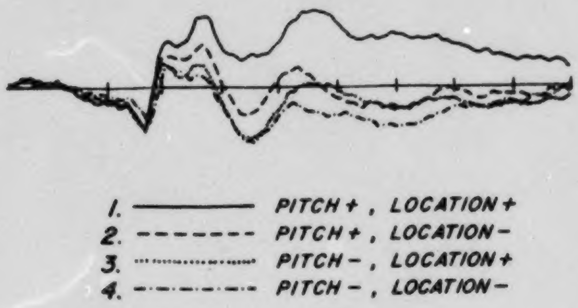

DIFFERENCE WAVES (F_z)

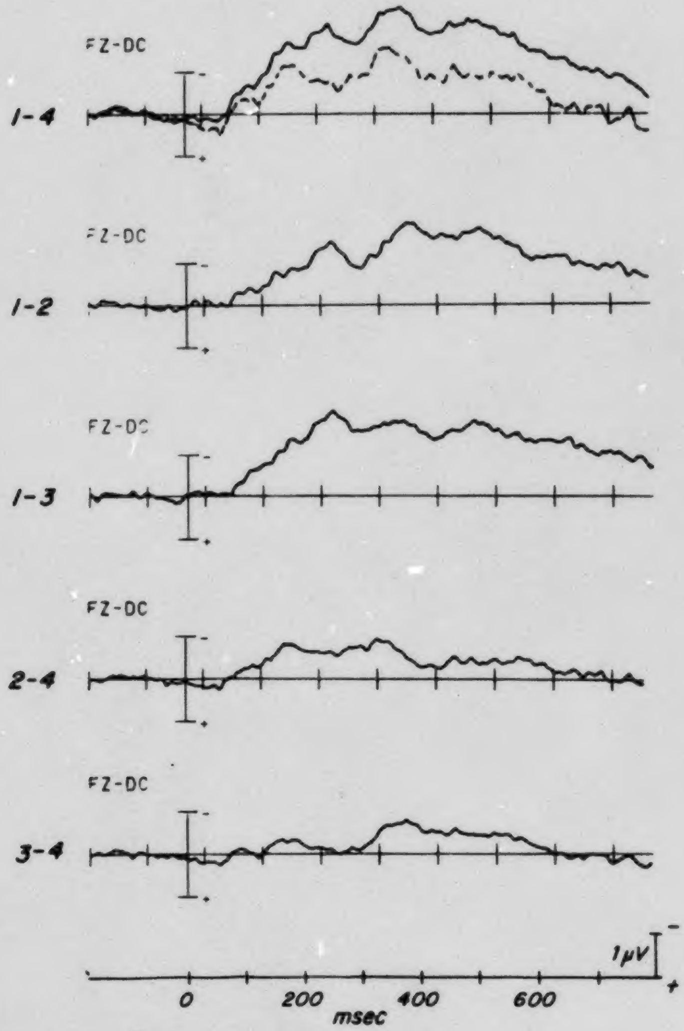

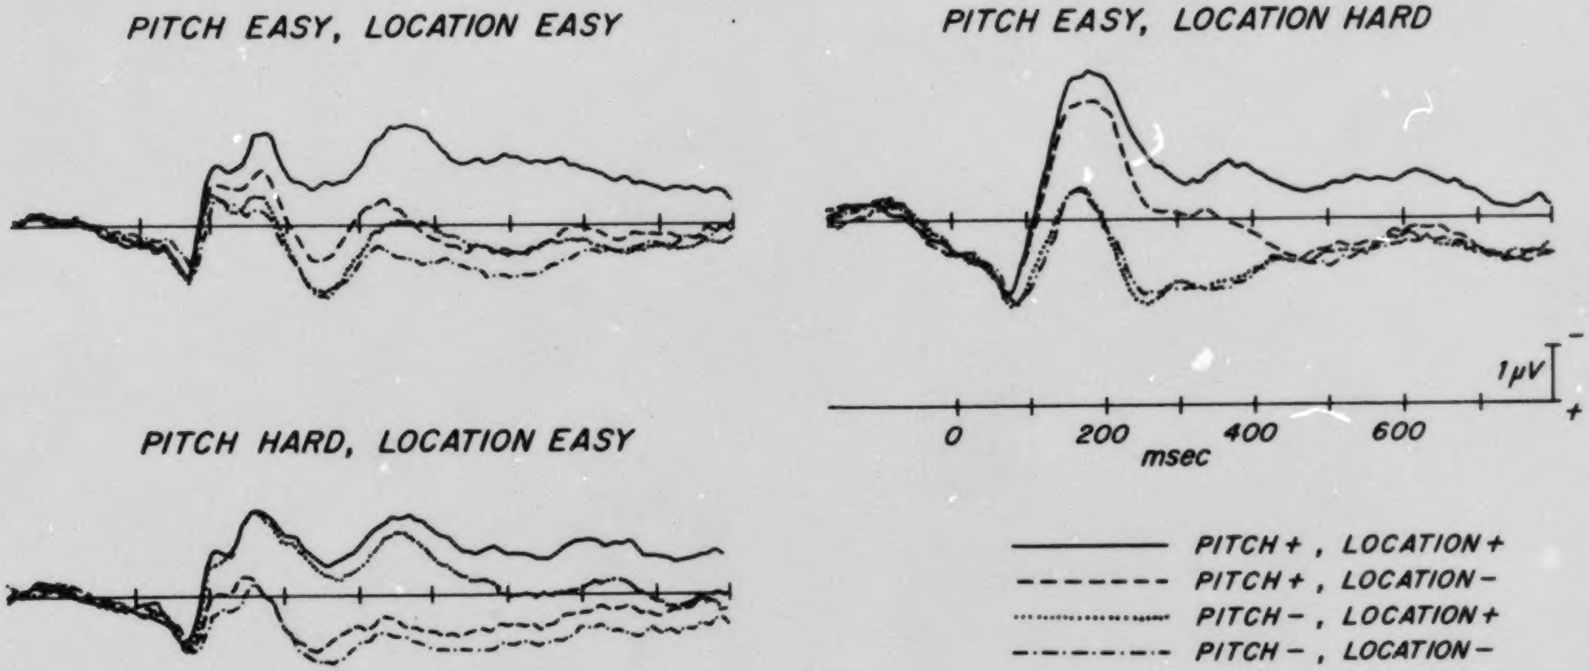
PITCH EASY, LOCATION EASY
PITCH EASY, LOCATION HARD
PITCH HARD, LOCATION EASY
1μV
−
+
0
200
400
600
msec
PITCH + , LOCATION +
PITCH + , LOCATION −
PITCH − , LOCATION +
PITCH − , LOCATION −

204

203

Remarks on Attention and Automaticity

Daniel Kahneman and Anne Treisman
University of British Columbia

We shall be discussing three topics, with which we have recently been concerned, jointly as well as separately. First, we cast a backward glance toward the early studies of attention, in an attempt to determine whether the newer theoretical approaches to this domain can accomodate the observations that led to the formulation of earlier theories. Our conclusion is skeptical. We then turn to some data recently obtained in our laboratory, which suggest that some prototypically "automatic" activities in fact depend on the allocation of attention, in a manner that casts some doubt on their automaticity. In the third part of this talk we question the assumption that perception varies along a single quantitative dimension of processing adequacy, depth or extent. Perhaps, as several authors have recently suggested, the functions of attention are better described by a clerical or bureaucratic model. If this view is accepted, many current problems, including the issue of the locus of selection and the nature of automaticity, may have to be reconsidered.

We have attempted to construct a list of major trends in the study of attention in the last decade. We have grouped these changes in three families: changes in the popularity of research designs, of concepts, and of theoretical notions. The direction of each change, toward increased or decreased popularity, is indicated.

TRENDS IN THE STUDY OF ATTENTION

•	
	CONCEPTS:
DOWN	Selection
DOWN	Filtering
UP	Automaticity
UP	Expectancy and Priming
•	
	TECHNIQUES
DOWN	Continuous tasks
DOWN	Selective listening
DOWN	Multi-stage tasks
DOWN	Accuracy measures
UP	Discrete trials
UP	Search tasks
UP	Expectancy costs and benefits
UP	RT measures
•	
	THEORY
UP	Expectancy theories of attention
UP	Mental life located in LTM

The changes of design have been particularly dramatic. The study of attention in the late 50's and early 60's was dominated by continuous tasks that required sustained attention. This was true in the early studies of vigilance, as well as in the studies of shadowing in selective listening. The stimuli were commonly presented in the auditory modality. The tasks had a relatively complex structure, and the quality of performance was measured by its accuracy.

All this has changed. The standard study of attention today consists of discrete trials, typically with visual stimuli. The tasks involve a small vocabulary of responses, as in search or detection. The favored measure of performance is reaction-time. The main manipulation of attention controls set and expectancy, rather than selective instructions. Many studies are concerned in one way or another with measurements of the benefits of confirmed expectations and of the costs of unconfirmed ones.

Turning to concepts, we note a marked decrement in the use of the notion of selective attention, and of the concept of filtering, both as a description of a particular research design and as the name of a possible mental operation. In contrast, we note the recent surge of interest in the notion of automaticity, as well as in the concepts of expectancy and priming as mechanisms of attention.

An obvious change in the nature of theorizing about attention is the increasing popularity of models that describe attention to an object as a state of anticipation. In its modern form, such a model of attention was first articulated by LaBerge (1973), and it plays an important role in the views of Posner (1978) and Shiffrin (1975; Shiffrin & Schneider, 1977). No one could deny the role of sets and expectations as aspects of attention, but the models of LaBerge and Shiffrin appear to treat anticipation as the sole mechanism of selective attention. The main weakness of such a theory is that it cannot explain filtering. Consider those ancient studies of selective listening, in which the subject was instructed to shadow a message presented to the right ear, and to ignore a message simultaneously presented to the left ear. The success of subjects in this task, even when the messages consist of randomly chosen words, cannot be attributed to biased anticipations, since the listener does not know in advance which words will be included in the relevant message, and is therefore incapable of preparing to recognize those words, or to ignore the words that will be heard on the irrelevant message.

The changes in theoretical attitudes and research practice shown in our list are highly coherent. The change of theoretical preferences in favor of expectancy models of selective attention was associated with a compatible change of research techniques: away from auditory messages, large vocabularies, and accuracy measures, toward discrete visual displays, small vocabularies, and reaction times. The result of these concurrent changes is that the dominant theoretical models of today explain the results obtained in the dominant research paradigms of today; however, our impression is that they no longer account for the results that had led earlier to the formulation of filter theories.

The loss of interest in the filtering paradigm has been associated with a shift toward late-selection models of attention, and with increasing interest in the notion of automaticity. The link between these developments is that late-selection models typically assume that the activation of logogens by their appropriate stimuli is automatic, effortless, and unconstrained by mutual interference at the central level (Deutsch & Deutsch, 1963; Keele, 1973; Shiffrin, 1975).

For reasons that will be developed below, the expectancy model of attention and the emphasis on automaticity are highly compatible with a view in which mental life is represented as a succession of states of long-term memory. This view has perhaps been articulated most clearly by Shiffrin & Schneider (1977). We call it a display board model of the mind. Imagine the mind as a board, in which each point is a light bulb that can be turned on, perhaps at several different brightness levels. The bulbs correspond to nodes in LTM, and they are connected in such fashion that activation may spread along some predetermined paths once a particular bulb is turned on. The lights on the display board correspond to dictionary units, or logogens. When our mental life is somehow concerned with a particular word, or another such unit of perception or thought, the light that corresponds to it will be turned on. In short, the lights are the devices by which we represent the identity of objects. A light can be turned on by the presence of its proper object in the stimulus field, so that the board serves the function of those displays that tell you who is in the office at a particular instant. In a different way (the difference must be indicated, or else we could not tell fact from fantasy), the lights are activated when we think about an object, rehearse its name as part of a list, etc.

The display-board model of the mind is quite powerful, and it provides an elegant account of many familiar effects. Let us first consider a partial list of phenomena with which it deals very well, before turning to other effects with which it deals rather poorly. A display-board model does well in explaining expectancy and priming effects. An expectancy is represented by a light bulb that is already preactivated; this is done in many electrical circuits, especially with vacuum-tube technology. When the tube is kept "warm", a small impulse will suffice to turn it on. The display-board model is also good for explaining spreading activation, by invoking the associative connections among nodes. If we assume that some connections are inhibitory, the model could also accomodate the costs of expectations as well as their benefits: a bulb could be selectively de-activated (by a competing expectation) so that a larger-than-normal signal would be required to turn it on. Thus, the display-board model is an excellent device to represent automatic activation. Indeed, it seems to have been designed for that purpose.

Without further elaboration, the board that we have described will not be able to simulate successful performance of a filtering task, since its only selective mechanism is a bias of expectation, which requires advance knowledge of the relevant and irrelevant stimuli. Another difficulty for the representation of filtering is that the display-board model (and the logogen system or long-term memory of most current theories) is organized in terms of responses. The various units of the board are accessed or activated by complex conjunctions of properties, but they map one-to-one on responses that the subject may be required to make. In a system of this type, it is simpler to designate a unit by the response to which it corresponds uniquely, rather than by any stimulus properties. In particular, it would seem natural for attention to be directed to classes of units that belong to a given domain of responses, such as digits, animal names or French words. Contrary to this expectation, one of the best-established facts of attention research is the superiority of stimulus set over response set (Broadbent, 1970; Kahneman, 1973). Later in this report, we consider alternative representations of mental activity, which can accomodate this observation.

The Problem of Automaticity

There appear to be two classes of claims about automaticity, which can be simply described as strong and weak. The strong version is associated with a definition of automaticity which was perhaps put most clearly by LaBerge (1975), and appears in a slightly different form in Shiffrin's work (Shiffrin & Schneider, 1977). It defines a process as

211

automatic if its performance is not facilitated by the allocation of attention. Weaker claims are associated with a weaker definition, which states simply that a process is automatic to the extent that it occurs without attention. The difference, of course, is that a process can be partly automatic in the weak sense, but still be boosted by the allocation of attention; in the strong sense of automaticity, a partly automatic process is not automatic.

The fact that many significant cognitive activities are (at least partly) automatic in the weak sense is not controversial. People may disagree on the degree to which this or that mental activity is automatic, or on the importance of the fact that it is automatic to this or that degree, but the weak automaticity claim could be on the platform of all political parties. Indeed, the notion of a filter which attenuates, rather than blocks rejected messages is equivalent to the assumption that semantic processing is partly automatic.

The feasibility of automatic processing, then is not at issue, although there may be discussion concerning the importance of attention in facilitating different classes of mental activity. The strong claim of automaticity, however, presents a sharp experimental question. Indeed, the strong claim has the distinction of being refutable. Refutability is perhaps overrated, and a psychologist who asks a colleague to state a position that can be refuted usually does so in the spirit of asking someone to stand still so that one can punch him on the nose. On the issue of automaticity, some theorists are obligingly standing still, and we shall take advantage of this fact.

Filtering in the Stroop Design

Reading familiar words is often invoked as the very prototype of a highly automatized skill. And the Stroop task is often invoked as an example of the automaticity of reading. The subjects are trying to identify the color of the ink in which a word is printed, but meanwhile the shape of the word automatically and speedily activates the corresponding logogen, thus causing interference. This appears to be the standard account of the Stroop task, and of the link between that task and the notion of automatic activation in reading, because of the demonstration that subjects read uncontrollably, even when it is to their best interest not to do so. We now describe some studies that question this interpretation of the Stroop effect, and the supposed automaticity of reading that the effect is said to illustrate.

Imagine a display that is tachistoscopically presented. The display consists of a square and a circle, which appear unpredictably on either side of the fixation point. The words RED and GREEN are printed, respectively in the circle and in the square. The word RED is printed in green ink, and the word GREEN is printed in red ink. Now imagine a display which is similar in all respects to the one just described, except that the words RED and GREEN exchange places, so that the color in which each word is printed corresponds to the meaning of that word. Consider a subject who is assigned the task of naming, as quickly as possible, the color of the ink in the circle. The correct response is 'Red' in both cases. Will it be made more easily and quickly in one of these cards than in the other?

The answer that one receives depends very much on who one is talking to. A lay person just laughs, because the result is intuitively obvious. Not so the attention theorist. Indeed, it is not at all obvious how a theory that contains the strong version of automaticity can explain a difference between the two conditions. Note that the subject is assumed to be fixating at the center, so that the quality of the sensory inputs is the same for both cards. If reading is automatic, then the logogens for 'Red' and 'Green' must both be activated by the printed words, equally on the two trials. Any facilitation or interference which is produced by such automatic activation should also be the same.

Several versions of this experiment have been run (Kahneman & Henik, 1981). The results clearly favor the lay person over the attention theorist. In one of the experiments, neutral words were used, as well as color words. The results are given below. The conditions are identified by capitalizing the word to whose color subjects were to respond. The words in the display could be neutral, compatible with the correct response, or conflicting. The results shown are mean correct RT in msec, and percent errors.

NEUTRAL-neutral	906	(3%)
NEUTRAL-compatible	944	(4%)
NEUTRAL-conflicting	956	(2%)
COMPATIBLE-neutral	858	(1%)
CONFLICTING-neutral	1108	(15%)

The results show significant interference by a color word (even a compatible one, in this case), which is distant from the area to which attention is directed. However, this effect is quite small in comparison to the effect of an incompatible word which is physically conjoined with the relevant color.

These results lend themselves readily to interpretation as an example of filtering in a discrete task. As in any other instance of filtering, several stages of processing are involved. The relevant circle is found at an early stage. Attention is paid to it. The allocation of attention to the circle facilitates all the responses that are associated with various aspects of that object. In particular, it facilitates the responses that belong to the set of color names, because these responses have been primed by the color-naming instruction. Thus, there appears to be no control over the activation of the logogen by the shape of the attended word, and in this sense the reading of the word is automatic. This automatic process, however, must depend on the allocation of attention, since the word in the square produces much less interference. It seems appropriate to ask how automatic an automatic process is, if it depends on prior selection? These results present an embarassment to theories that claim reading to be automatic in the strong sense. A similar conclusion has been reached by Francolini & Egeth (1980), on the basis of rather similar experiments.

The interpretation that we suggest assumes that visual attention is especially effective when it selects an input (Treisman, 1969) or an object (Kahneman, 1973). When an object has been selected, another selective operation must be invoked to determine which property of the object will be allowed to control responses. In general, the priming of a response category is enough to do most of the work of selection, because different properties of an object are rarely linked to different members of the same class of responses. The Stroop task is an exception, of course. It produces interference because an irrelevant property of the selected object is strongly associated with a response that has been primed by the task. The visual suffix effect, and perhaps the auditory suffix effect as well, could be interpreted in the same manner: an irrelevant member of a relevant group of items is not easily distinguished from its relevant neighbors, and causes as much interference with their processing as an extra relevant item would do (Kahneman, 1973; Kahneman & Henik, 1977, 1981).

Visual filtering is a robust effect, which we have demonstrated in several experiments. In one of the studies in this series we presented two words, on either side of the fixation point. One of the words was always printed entirely in black, the other was printed entirely in colored ink, or had a single colored letter in it. The subject's task was to report the color of the ink. Here again, Stroop interference was much more pronounced when the colored ink and an incompatible color name were conjoined than when they were spatially separated (Kahneman & Henik, 1981). An interesting result was obtained when a single letter was colored. We had expected the control RT to be longer in this case than when a whole neutral word was printed in color, simply because a single letter is more difficult to find than a whole colored word. The expected difference was found, and it was significant. In addition, we predicted that naming the color of a single colored letter embedded in a color name would be associated with less interference than naming the color of the entire word. Focused attention to a single letter would reduce the probability, or the speed, of reading the incompatible word. Indeed, the measure of Stroop interference was 73 msec for the letter and 159 msec for the whole word.

The results of these experiments illustrate the concept of filtering in visual presentation, and present some difficulties for the strong claim concerning the automaticity of reading, and for the interpretation of Stroop interference as evidence of such automaticity. The major conclusion is that it is essential to distinguish selection of inputs, or objects, from selection of properties. As we have seen, observers are capable of efficient rejection of irrelevant objects, but the irrelevant properties (and perhaps parts) of a relevant object cannot be prevented from contacting their traces, and from activating irrelevant responses. The distinction between selection of objects (or inputs) and selection of properties (or analyzers, in Treisman, 1969) seems salient and fundamental, yet it is often ignored in psychological research and theory. The difference between objects and properties was lost historically when the ambiguous term "stimulus" was adopted in psychological theorizing, because it can be said either that the object (e.g., a red O) is the stimulus, or that a property (redness or circularity) is the stimulus. In the behaviorist tradition, the term was applied to "whatever controls a response". Because discriminated responses are controlled by properties (peck the key if the cage is illuminated, not if it is dark), it is most natural in that tradition to think of stimuli primarily in terms of properties, and to ignore the notion of objects altogether. This legacy has influenced current studies of information processing. It is illustrated by treatments that interpret Stroop interference as a failure of selective attention and as evidence of the automaticity of processing. In fact, the Stroop effect demonstrates that people do not easily select among the properties of the attended object. On the other hand, the present results add to the extensive evidence from other studies of filtering to support the conclusion that irrelevant objects can be rejected quite effectively.

The Dilution Effect

We now turn to another set of tests of the notion of automatic activation, which we label the dilution effect. These tests are addressed to a straightforward prediction of the strong version of automaticity, as formulated by LaBerge(1975). If automatic activation of a node or logogen by a familiar stimulus does not require attention, then the activation of that node should not be affected by the presence of other objects in the field.

A strong claim of automatic processing could be supported in either of two ways. The first is a demonstration that the efficiency of processing a relevant stimulus is unaffected by advance knowledge of its nature or location. The second is a demonstration that efficiency is also unaffected by the concurrent presentation of irrelevant stimuli. The first procedure was introduced in LaBerge's classic experiment, which attempted to show that the processing of familiar stimuli is independent of attention, by demonstrating that such stimuli are processed no faster when they are specfically expected than when they are not, in contrast to unfamiliar stimuli, which show benefits of expectations. A well-known series of studies by Shiffrin and his collaborators (summarized in Shiffrin, 1975) appeared to make much the same point, by showing that prior knowledge which permitted the focusing of attention on a channel sometimes failed to facilitate the processing of a stimulus presented on that channel, when other stimuli were presented at the same time. More recent experiments (Duncan, 1980; Posner,

1978) using the methods of Shiffrin and LaBerge respectively, have raised some doubts about the absence of anticipation effects in these designs. However, the absence of display size effects in some search tasks provides incontrovertible evidence for the automaticity of that operation, at least under some conditions (Egeth, Jonides & Wall, 1972; Schneider & Shiffrin, 1977; Treisman & Gelade, 1980).

If a search task can be performed in parallel over an entire display, then the performance of that task is automatic in the strong sense. A filter theorist would say that such search tasks are performed by pre-attentive mechanisms, since parallel operation is the defining attribute of these mechanisms: Broadbent has commented on the fact that his S-system and P-system had been unfortunately labeled, since the operations of the former system are parallel and those of the latter are serial. Schneider & Shiffrin (1977) established an important fact that would not have been anticipated in filter theory: with prolonged practice and consistent mapping of stimuli to responses, search could be automatized even when the targets shared no obvious physical features. It is important to note, however, that the response that was automatized in these experiments was a simple instrumental response. There is no evidence in this work that automatization can occur when a large set of stimuli is mapped, even consistently, onto a large set of responses.

The distinction between filtering tasks and search tasks is defined by the nature of the mapping of stimuli to responses.

The tasks of search and detection radically simplify the problem of response selection, since the response is completely determined as soon as the relevant stimulus has been detected. In a filtering task, such as selective shadowing, or report of the middle row of a Sperling array, a problem of response selection remains after the relevant part of the message has been found. It is not enough to find the relevant row, the particular sequence of digits must be produced. Reading is the prototypical instance of an information-preserving activity, in which the richness of the stimulus ensemble is preserved in the repertoire of responses. It is therefore of interest to ask whether the "automatic" reading which is assumed to occur, for example, in the Stroop situation is subject to an effect of display size. The strong claim of automaticity allows no mechanism that could produce such an effect, except sensory interference.

We have studied this question in a long series of experiments. In a typical study, the subjects are shown a colored bar centered on the fixation point, and are asked to name the color of the bar as quickly as possible. On some trials, a single word is presented, some distance above or below the bar. The word is sometimes unrelated to the color-naming task, sometimes compatible with the response that is correct on the trial, and sometimes is the name of another color. As many other investigators have found (Dyer, 1973; Gatti & Egeth, 1978), the presentation of a color name affects the speed with which the color of the bar is named. Interference and

221

facilitation are both obtained, although the magnitude of the effects is smaller than when the relevant color and the irrelevant word are conjoined. The occurrence of Stroop interference in this situation represents at least a partial failure of selective attention to objects. The relevant color bar is presented at the fixation point and the subject is encouraged to focus attention on that area; the word is irrelevant and its reading is surely involuntary. The purpose of the present set of experiments is to determine whether the reading of irrelevant words is also 'automatic', in the sense defined by the absence of a display-size effect. We introduce a minimal variation of display size: on some trials, an extra word is presented, on the other side of the relevant color bar. Our question is whether the presentation of the added neutral word will affect the amount of interference or facilitation that the color word produces.

We have run this experiment, in diverse variants, with more than 100 subjects. The consistent result of these studies is that the facilitating effect of a compatible color name and the interfering effect of an incompatible one are both reduced by the presentation of a neutral word. We call this result the dilution effect, since the impact of a task-related, distracting word is reduced by the presentation of another word, which is unrelated to color naming. In several experiments we have observed that adding a neutral word to the display retards color-naming in the presence of a compatible color word (by reducing the normal faciliation which that word would produce) and improves the speed of color- naming, in the presence of a color word which is incompatible with the correct response. The effects are small (interference may be reduced from 90 to 60 msec, and facilitation from 60 to 40 msec), but they are highly consistent.

We have compared the dilution effect which is produced by adding an extra word to the display to the effect of adding a row of X's. Much to our surprise, the effects turned out to be nearly identical. In one experiment, a word diluted the interfering effect of a conflicting color name by 20 msec, and the facilitating effect of a compatible one by 25 msec. The corresponding values for a row of X's were 19 and 15 msec, and the difference between the conditions was not significant. We have replicated this result ad nauseam. It is solid in the context of the Stroop experiment, although we have found other situations in which a neutral word causes greater dilution than a row of X's.

The dilution effect presents an embarassment to the claim that involuntary reading is automatic and unaffected by attention. (We suppose that no one would wish to extend that claim to voluntary reading). If the visually presented words automatically access their corresponding nodes in long-term memory, and if this access automatically activates primed responses, the number of words that are simultaneously presented in the display should have no effect on the efficacy of this priming. However, the validity of our argument depends on our ability to demonstrate that the two items do not interact at a peripheral level. This turns out to be quite difficult, because the notion of peripheral interaction has recently been broadened to include various types of feature interactions and perturbations (Estes, 1972; Wolford, 1975, 1980), and the conditions for these interactions are only poorly specified. It seems reasonable to argue that one of the defining properties of peripheral interactions is that their magnitude should depend on the retinal distance between the interacting elements. In several experiments, however, we have found that the distance between the diluting word (or row of X's) and the color name had little effect on the facilitation or interference produced by that name.

We conclude that the dilution effect provides evidence against the notion that involuntary reading is automatic, effortless and independent of attention. However it is important to note that the notion of automatic reading assumes the automaticity of two distinct processes: (i)the activation of the appropriate node; (ii)the spreading of activation from this node to other nodes, or to structures that control responses. The dilution effect suggests that one of these events is not automatic, and is subject to interference by the concurrent processing of other stimuli. What is the nature of this dilution effect?

The following possibilities come to mind:

(i)The dilution effects are due to peripheral interactions that reduce the signal-to-noise ratio in the sensory message. In this view, dilution is an artifact. We have mentioned some of

224

our reasons for believing that it is not.

(ii)The processing of information is facilitated by the allocation of some resource to the relevant region. Dilution occurs when these limited resources are distributed among several regions or objects of processing.

(iii)The onset of a stimulus elicits an automatic orienting response. (Posner, 1978; Posner, Nissen & Ogden, 1978). The presence of several stimuli in the field induces a conflict among competing orientation tendencies, and fully effective processing is delayed until this conflict is resolved.

(iv)The accumulation of information in recognition nodes (or logogens) is not affected by the concurrent accumulation of information in other nodes, but the efficacy of outputs from any one node in controlling responses or in activating other nodes is reduced by concurrent activity elsewhere in the system.

(v)Simultaneously present stimuli all evoke response tendencies and the appearance of a dilution effect is produced by response conflict (Eriksen & Eriksen, 1974).

To advance the choice among these possibilities, we have used a situation that we label the dummy-conflict design. In the basic variant of this design, the subject is shown either a single word, which appears unpredictably above or below the fixation point; or two words, one above and one below fixation. The instruction is to read either of the presented words, as quickly as possible. An obvious argument could be made that the response to the dual display can only be faster than the response to a single word, since the more rapidly processed member of the pair should control the response. In fact, the presentation of two words causes a slight, but highly consistent slowing of the reading time: from 521 to 555 msec in a typical experiment. Here again, the distance of the words from the fixation point does not seem to matter very much. An interference effect of 25 msec was still found when the nearest contours of the interacting words were 3 degrees apart, across the fovea. Furthermore, the effect did not depend on the presentation of a second readable word: a word-sized patch of random dots on one side of the fixation point caused about as much interference as an additional word.

It is of course not certain that the interference observed in the present experiment has the same origin as the dilution effect which was discussed previously, but the hypothesis that the two effects are similar is both plausible and parsimonious. The observations made in the dummy-conflict situation permit the elimination of some of the possible interpretations that were mentioned earlier in the context of the dilution effect. Hypothesis (i), which attributes the interference to peripheral interaction, is not supported by its insensitivity to retinal distance. Hypothesis (v), which attributes the interference to conflict between incompatible responses, cannot explain the effects of the dots. Hypothesis (iv), which assumes interference between the outputs of different logogens is also rejected, for the same reason. The remaining hypotheses attribute the interference to competition for limited processing resources or to a conflict between orienting tendencies, which

226

must be resolved before the processing of a particular region of the field can be carried out at full efficiency. We hope to test these possibilities in future experiments.

Concluding Remarks

We believe that our experimental results provide substantial evidence against the position that we have labeled the strong claim of automaticity. We found that the involuntary reading of Stroop words was affected by the allocation of attention to a spatial position, and also by the presence of another word, or row of X's. The accuracy and the speed of voluntary reading are also affected by the presence of an irrelevant item in the field, even when the distance between words makes peripheral interaction unlikely. In our experiments, then, reading does not appear to be fully automatic. This conclusion, however, appears to be in conflict with recent evidence of the potent effects of words which are 'read' without awareness (Marcel, 1981). In a provocative series of experiments, Marcel observed that words that are masked by pattern, to the point that observers do not know they are there, nevertheless can prime responses to a subsequent target in a lexical decision task, and cause interference or facilitation in the naming of a color patch.

In our laboratory, Peter Forster has confirmed a highly reliable (although frustratingly small) facilitating effect of a masked color word on the naming of a subsequent color patch. We have both been subjects in the experiment, and both showed a substantial effect of words that we never consciously 'saw', although we knew they were shown on some trials and were watching for them. Evidently, there was some reading of the unseen words, and it appears very natural to apply the label 'automatic' to this reading. How can this conclusion be reconciled with our other results?

It seems easy to achieve a superficial reconciliation. One possibility is that the activation of logogens by subliminal stimuli occurs because reading is partially automatic. In other words, we can read without consciousness, but perhaps at a lower grade level than usual! It may be a mistake to equate attention with conscious experience. In Marcel's subliminal experiments the subjects were certainly attending to the location of the 'unseen' words. It is possible that attention limits could also be shown subliminally; for example, interference and facilitation from subliminal words might be diluted by the addition of subliminal neutral words.

There are some indications that such patching of our models may not be adequate. A more fundamental revision will be required, if evidence is obtained for a pattern that we label strong dissociation. Marcel's results are an instance of a dissociation between phenomena that we expect to be linked. They reinforce the conclusions of a large literature on subliminal effects, confirming dissociations between awareness and other indications that a message has been registered and understood (Dixon, 1971). These results are often compatible with the notion that different response systems apply different criteria to the *same* information. In strong dissociation, however, unconscious processing is as precise and refined as conscious processing would be, and it is therefore not possible to argue that available information did not meet the higher criteria required for registration in consciousness, although it sufficed to trigger some low-level responses. The conclusion from strong dissociation would be that information which is available to control some responses is not made available to conscious elaboration. The evidence for strong dissociation is very scanty, but intriguing. Thus, von Wright, Anderson & Stenman (1975) reported that the emotional responses to synonyms of shock-associated words (but not to the words themselves) were the same regardless of whether these words were presented on the attended or on the unattended channel, and Marcel (1981) reported some data in which the priming effects of subliminal and supraliminal stimuli were of the same magnitude. Confirmation and extension of such findings would force many of us to reconsider our presuppositions about information processing and conscious experience.

Dissociation phenomena are troubling because they raise doubts about the validity of our sense of personal identity. If my skin responds to the emotional significance of words that I have not seen, do *I* know the meaning of these words? There seems to be no good answer to this question. The solution to the dilemma may be to revise our criteria for the use of epistemic words such as "know", "see", or "understand". In particular, the suggestion has been made by authors in the traditions of artificial intelligence (Minsky, 1975), philosophy (Dennett, 1978) and experimental psychology (Allport, 1980) that we treat the mind as a collectivity of semi-independent entities, rather than as a single one.

Perhaps we should take as our model of the mind a large organization, such as General Motors, or the CIA. Under what conditions can such an organization be said to know something? Certainly, the organization "knows" a fact if all the significant individuals within it know it, as shown by their actions. But there are many borderline cases. Does the CIA know a fact if one functionary in that organization knows it, but has told nobody else, or is believed by no one? Does an organization know a fact if the lower echelons act on it, but without informing higher echelons that they do so? The observation of strong dissociation phenomena suggests that it may be as difficult to assign epistemic states to individuals as it is to assign such states to organizations.

It now appears at least conceivable that future discussions of attention may be conducted within the framework of an organizational metaphor for the mind. It is disconcerting, but perhaps also encouraging, that many of the questions with which we have been concerned for years --including the question of automaticity that was the focus of this paper-- will turn out, in such a framework, to be slightly out-of-focus. The proponents and opponents of the idea of automatic semantic processing, just like the proponents and opponents of early selection, shared many presuppositions. In particular, they shared the idea of a standard path of information processing, and the idea that attention operates at one or more bottlenecks (or road-blocks) along that path, to select the messages that should be processed further, or perhaps to attach to them a single value of relevance. While we continue the debate within the old framework, we should remain alert to the possibility that it could soon become obsolete.

REFERENCES

Allport, D.A. Patterns and actions: Cognitive mechanisms are content-specific. In G. Claxton (Ed.), *Cognitive Psychology: New Directions*. London: Routledge & Kegan, Paul, 1980.

Broadbent, D.E. Stimulus set and response set: Two kinds of selective attention. In D.I. Mostofsky (Ed.), *Attention: Contemporary Theories and Analysis*. New York: Appleton Century Crofts, 1970.

Dennett, D.C. *Brainstorms*. Ann Arbor: Bradford Books, 1978

Deutsch, J.A., & Deutsch, D. Attention: Some theoretical considerations. *Psychological Review*, 1963, *70*, 80-90.

Dixon, N.F. *Subliminal Perception: The Nature of a Controversy*. London: McGraw Hill, 1971.

Duncan, J. The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, 1980, *87*, 272-300.

Dyer, F.N. The Stroop phenomenon and its use in the study of perceptual, cognitive and response processes. *Memory and Cognition*, 1973, *1*, 106-120.

Egeth, H., Jonides, J., & Wall, S. Parallel processing of multi-element displays. *Cognitive Psychology*, 1972, *12*, 278-286.

Eriksen, B.A., & Eriksen, C.W. Effects of noise letters upon the identification of a target in a nonsearch task. *Perception and Psychophysics*, 1974, *16*, 143-149.

Estes, W.K. Interaction of signal and background variables in visual processing. Perception and Psychophysics, 1972, 12, 278-286.

Francolini, C.M. & Egeth, H.E. On the nonautomaticity of 'automatic' activation: Evidence of selective seeing. Perception and Psychophysics,1980,27, 331-342.

Gatti, S.W. & Egeth, H.E. Failure of spatial selectivity in vision. Bulletin of the Psychonomic Society, 1978, 11, 181-184.

Kahneman, D. Attention and Effort. Englewood Cliffs, N.J.: Prentice-Hall, 1973.

Kahneman,D. & Henik, A. Effects of visual grouping on immediate recall and selective attention. In S. Dornic (Ed.), Attention and Performance VI. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Kahneman, D., & Henik, A. Perceptual organization in attention. In M.Kubovy & J. Pomerantz (Eds.), Perceptual Organization. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

Keele, S.W. Attention and Human Performance. Pacific Palisades, California: Goodyear, 1973.

LaBerge, D. Attention and the measurement of perceptual learning. Memory and Cognition,1973,1,268-276.

LaBerge, D. Acquisition of automatic processing in perceptual and associative learning. In P.M.A. Rabbitt & S. Dornic (Eds.), Attention and Performance V. London: Academic Press, 1975.

Marcel, A.J. Conscious and unconscious perception: The effects of visual masking on word processing. *Cognitive Psychology*, 1981, in press.

Minsky, M. A framework for representing knowledge. In P.H. Winston (Ed.), *The Psychology of Computer Vision*. New York: McGraw Hill, 1975.

Posner, M.I. *Chronometric Explorations of Mind*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.

Posner, M.I., Nissen, M.J., & Ogden, W.C. Attended and unattended processing modes: The role of set for spatial location. In H.L. Pick (Ed.), *Modes of perceiving and processing information*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.

Schneider, W. & Shiffrin, R.M. C Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review*, 1977, *84*, 1-66.

Shiffrin, R.M. The locus and role of attention in memory systems. In P.M.A. Rabbitt & S. Dornic (Eds.), *Attention and Performance V*, . London: Academic Press, 1975.

Shiffrin, R.M. & Schneider, W. Controlled and automatic human information processing. II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 1977, *84*, 127-190.

Treisman, A. Strategies and models of selective attention. *Psychological Review*, 1969, *76*, 282-299.

Treisman, A. & Gelade, G. A feature integration theory of attention. *Cognitive Psychology*, 1980, *12*, 97-136.

Von Wright, J.M., Anderson, K. & Stenman,U.

Generalization of conditioned G.S.Rs in dichotic listening. In P.M.A. Rabbitt & S. Dornic (Eds.), *Attention and Performance V*. London: Academic Press, 1975.

Wolford, G. Perturbation model for letter identification. *Psychological Review*, 1975, *82*, 184-199.

Wolford, G. & Shum, K.H. Evidence for feature perturbations. *Perception and Psychophysics*, 1980, *27*, 409-420.

THE ROLE OF ATTENTION IN OBJECT PERCEPTION

Anne Treisman
University of British Columbia

I'd like to suggest a new view of attention — or at least of one role it may play in perception. Let me set the scene by flashing the first slide briefly. You probably all had an immediate impression of recognizable objects, organized coherently in a meaningful framework; a familiar experience. Analysis into more elementary sensations is difficult and feels abnormal; yet in order to perceive and identify any complex object, we must normally register not only its features (shape, size, color, etc.), but also the fact that they are conjoined in a particular configuration. We see a rose as red, its leaves as green and the vase as grey. This allocation of colors to shapes, sizes to orientations, voices to locations, seems immediate and automatic.

However, there has been a long controversy in Psychology about how we achieve this impression of unitary objects. Do we build the world we perceive out of more elementary sensations, or do we directly register wholes, meanings, relations? For a time the Associationist or building-block view was generally accepted; complex objects were painstakingly pieced together from more elementary sensations. But it has alternated with strong swings to the other extreme, for example with the Gestalt movement, and now again it is more fashionable to talk of 'top-down' processing, and of 'frames', and to emphasize the wholistic quality of immediate perception. Intuition and introspection push us compellingly towards the Gestalt position. We do see people, plants, buildings; not colors, edges, and movements. However, the immediacy and directness of a conscious experience are no gurantee that it reflects an early stage of information processing in the nervous system. We have little idea what the relation between subjective experience and neural coding could be. This claim can therefore only relate to the order in which we can introspectively access different levels of representation. We may become aware only of the final outcome of a complicated sequence of prior operations. 'Top-down' may describe our experience, but for a theory about perceptual coding, we need other kinds of evidence.

Recent physiological research shows that the early stages of coding depend on specialized populations of cells, tuned selectively to discriminate particular properties. These properties seem to be mapped into different areas of the brain (Zeki, 1976). If we look at Psychology, we also find evidence that at least some dimensions are 'separable' (Garner, 1970); they can be processed independently and in parallel, they can be selectively attended to, and they contribute independently to judgments of similarity. Even phenomenological evidence is not unequivocal: The next two slides will show that there is some difficulty in perceiving complex pictures when the context is quite unfamiliar. You may, in fact, have registered mainly colors and lines rather than meaningful objects. It seems at least possible that our normal glimpses out of the corner of the eye in everyday life are equally imprecise or ill-defined, without our being aware of it.

I would like, at least provisionally, to accept that there is some decomposition of the visual input into separable parts and properties. This at once raises the question how these properties get put together again, and how we normally avoid mistakes in the form of "illusory conjunctions". Why don't we see a gray rose with red leaves in a green vase? One answer is that in natural contexts there are many known constraints limiting which features we can sensibly combine: grass is normally green, the sky blue or grey, people have noses and tables do not. Perception maps sensory data into expected 'frames', selecting the combinations of features that make sense. On the other hand, in less predictable situations, a rapid glance or a diffuse and inattentive gaze may in fact generate some hallucinatory couplings, without our always taking the time to check their validity. A friend walking in a busy street 'saw' a colleague and was about to address him, when he realized that the black beard belonged to one passerby and the bald head and spectacles to another.

So how do we normally avoid making these 'conjunction errors' whenever the context leaves some ambiguity? I would like to suggest a new hypothesis about the role of attention, which links it directly to this problem of object perception. I will state the hypothesis in a fairly extreme form, because it makes it both clearer and easier to refute. But I would like, from the outset, to stress that it is still a tenative proposal, whose implications I am still actively testing and exploring.

My suggestion is that we specify which features should be integrated into one object by scanning each location serially with focal attention. Any features that we register within the same central 'fixation' are integrated into a single percept. In other words, focal attention provides the glue which puts objects together. I believe there are two distinct levels of processing; at the first, all the separate features are coded independently and in parallel to form different feature maps, using populations of feature detectors for colors, orientations, directions of movement, spatial frequencies. This stage organizes the perceptual world into groups, textures, homogeneous areas — candidates for possible objects and events to be identified at the next stage using focused attention. It is similar to the preattentive stage first proposed by Neisser in 1967, although for me it is also preconscious, and there are one or two other differences which I will discuss later.

Before any features can be consciously perceived, they must be recombined. Conscious experience is populated not by disembodied orienations, colors and movements, but by objects and their backgrounds. Feature registration is followed, then, by a feature-integration stage, where selective attention mediates the formation of multi-dimensional objects. Whenever the context is unfamiliar, so that we cannot use prior knowledge to match up the features of objects, an act of attention is necessary to ensure that correct conjunctions are formed. We attend to one object at a time; in other words, serial processing is necessary.

I had this idea four or five years ago and since then I have been trying to test it. Let us see what predictions it generates.

1) If an object has a unique feature, we should be able to detect or identify it independently of the number of other objects which are present. So search for a feature target should be parallel across a display.

2) If an object has no unique feature and is defined only by a conjunction of features, identifying it should require focal attention to each item in turn, and should therefore force serial processing.

3) It seems likely that texture segregation and figure-ground grouping are pre-attentive, parallel processes, they should therefore be determined only by features and not by conjunctions of features. At this pre-attentive level, the theory claims that there are many separate worlds — a world of color, a world of shape, a world of movement, all organized along their own

particular lines, but not related or unified until attention is brought to bear. Here is one major difference from Neisser's single, global preattentive world.

4) If attention is prevented or overloaded, features will be "free-floating" with respect to one another. This means that illusory conjunctions could be formed when more than one object is present, by wrongly recombining the features of different objects. For example, the red blouse and black hair of the woman in the first slide were conjoined by one subject to give a red-haired woman.

5) Without attention, features could also be *spatially* free-floating. In other words, we may not know exactly where they are located, although we can certainly home in on them very rapidly. This could not be the case for conjunctions. If we identify a conjunction, we must first have located it.

6) Finally, we should be able to use these phenomena to identify new or doubtful features. For example, is closure an elementary feature? If so, it should allow parallel search and texture segregation and it should join in illusory conjunctions when attention is overloaded.

I have tried to test these predictions in a number of different paradigms and I'd like to describe some of the results. The first experiments test the serial search prediction — the idea that attention must be directed to each item in turn whenever the conjunction of features is relevant. The next three slides illustrate the task. In the first two, the targets to look for are either the letter 'S' or a blue letter. In the third, the target is a green 'T'. You probably found that the target 'jumped out' of the background of green 'Xs' and brown 'Ts' in the first two cases and required a painstaking scan in the third case, unless you were lucky and happened to focus on it early. In the experiment, we tested two conditions of visual search. In the first, the target was defined by a conjunction of features (green 'T'), while in the second it was defined by two disjunctive features (either 'S' or blue). Subjects were to press one key if a target was present and the other key if no target was present. We measured the time they took. The distractors were always green 'Xs' and brown 'Ts'. Notice that subjects had to check two dimensions in both conditions, but needed to check a particular spatial combination only in the conjunction condition. If focal attention is required for conjunctions, the scan should be serial and self-terminating. We therefore predict (a) that search time will increase linearly with the number of items in the display; (b) that the positive slope should be half the negative slope, because on average, when the target is present, subjects will find it after checking half the items. For the feature targets, neither prediction should hold. Figure 1 shows the results. It shows a flat or non-monotonic relation between display size and detection time for the feature targets. There is no evidence that subjects scan serially in order to find the blue letter or the 'S', targets that are defined by disjunctive single features. When the target is present, it apparently jumps out at you; no scanning is necessary. When it is not present, subjects are much slower and do take longer, the more items are present. They may scan the display to make sure.

With conjunctions, on the other hand, we have a very different pattern. Both functions are straight lines, linearity accounting for 99.7% of variance due to display size. The ratio of positive to negative slopes is 0.43, which is quite close to one half. Notice that we cannot explain this different pattern of performance by the similarity or confusability of targets and distractors, because this should affect displays of one in the same way as the larger displays. The difficulty with conjunctions arises only when more than one item is present. So far, then, the results agree with the predictions.

How general is this finding? What happens with practice on conjunction search. Can we get rid of the serial scan and automatize the search for conjunctions? Can we perhaps set up a new unitary detector for a "green 'T'" which will allow it to jump out of the display without focal attention. We ran four of the subjects for seven sessions and two of them for thirteen sessions and found no evidence that the pattern was changing. The slope remained just as linear on block 13 as on block 1. The slope and intercept both decreased with practice but most of the change came in the first three blocks. It is possible that if we had doubled the number of sessions we might eventually have found a change, but there is little sign of one here. It is interesting to speculate that there may be built-in neural constraints on which properties can be unitized. Perhaps we just aren't built to respond automatically to green 'Ts'.

Let's look at one or two obvious variations of the task. Suppose we make the features easier or harder to discriminate. What should happen, if my hypothesis is correct? It should make each of the serial checks for the

240

conjunction target either slower or faster; so it should change the slope relating search time to display size without changing the pattern of linear functions and the two to one ratio of negative to positive slopes. That is what we found. We had subjects searching either for a red 'O' in green 'O' and red 'N' distractors, where discriminability was high both for shapes and for colors, or for a green 'T' in green 'X' and blue 'T' distractors, where discriminability at the feature level was considerably lower. The search rates (see Figure 2) vary quite dramatically (92 msec per item for confusable colors and shapes to 40msec per item for easily discriminable ones). Yet they preserve the pattern of serial, self-terminating search. This is important because it confirms that the difference between the conjunction and the feature conditions is not simply a difference in difficulty. We might have argued that conjunction search requires attention only because it is harder than feature search. That doesn't seem to be the case.

Another question: is the serial scan really an eye movement scan rather than an attention scan? Is it our mental or our physical eyeball that we move successively from item to item? If we use smaller displays more centered on the fovea, will we change the pattern? The next slides show two new conditions, one with densely packed items and one in which they are spread over an area which is twice as large. Figure 3 shows the results. The scanning rate is the same for the dense and the sparse displays. The serial self-terminating scan is clearly not a function of retinal distance, as it would be if visual acuity and eye movements were determining the search time. The results are more consistent with the idea that the serial fixations are made with a mental rather than a physical eyeball — with our internal attentional spotlight.

A final important question about generality: Do all these findings apply only to color and shape conjunctions, or are they generally true of separable features? For example, could they hold up with local elements of more complex shapes — lines, angles, curves, etc. Must we use focused attention to tell us, for example, whether the cross-bar in a particular display goes with the two diagonals to form an 'A' or with the two verticals to from an 'H'?

I tested this question by choosing sets of letters that could potentially give rise to illusory conjunctions, compared to letters that could not. One target in the conjunction condition was an 'R' in a background of 'Ps' and

'Qs'. The 'R' could be made out of the 'P' and the diagonal line of the 'Q'. Another was a 'T' in a background of 'Zs' and 'Is'. Would these targets therefore require a serial scan? Or is the difficulty of search through letter sets determined simply by how confusable the targets and distractors are? I used as control conditions, search for letters which, taken individually, were more confusable with the targets: 'R' in 'Ps' and 'Bs' and 'T' in 'Ys' and 'Is'. Figure 4 shows the results: Again, we have linear functions and apparently serial search through the letters that I hypothesized might be at risk for illusory conjunctions, and faster, non-linear functions for the others. The ratio of positives to negatives was close to 1/2 for the conjunctions and much less for the similarity controls (0.27). Again, we have evidence for two separate underlying processes for positive and negative decisions with the similar letters, and a single serial process for the conjunctions. The order of difficulty reverses when the distractors are homogeneous. Target 'R' in 'Bs' is then harder than in 'Qs'. It is only when we add 'Ps' to both that conjunction errors with 'Q's become possible and focused attention becomes necessary.

The conclusion seems to be that even with such higly familiar stimuli as letters, the risk of illusory conjunctions can arise and force the use of focal attention to integrate the local features into the correct units. I am not claiming that serial scanning is necessary for all sets of letters. Any that can be distinguished by single features (e.g., the presence of a curve or a diagonal) should potentially be recognized in parallel. I think this experiment has also shown that similarity between individual items is not the most powerful variable determining attention limits in visual processing. This is of course relevant to other models (for example, Gardner, 1970), which have claimed that perception is based on unlimited parallel channels, and that confusion errors arise only at the decision level.

The next prediction concerns texture segregation. Early detection of boundaries is a primary requirement in perception. Before we can identify an object, we must segregate it from its background. If texture segregation and figure-ground grouping are pre-attentive processes, depending on parallel registration across the visual field, the theory predicts that they should be determined only by separate features and not by changes in the conjunctions of features. So we should group two sets of items easily on the basis of color (e.g., red shapes vs. blue ones) or a simple feature of shape (e.g., curved vs. straight), but not on the basis of conjunctions of these properties, (red curved and blue straight letters vs. red straight and blue curved letters). The next three slides confirm these predictions.

The theory actually makes quite a strong and somewhat surprising claim about this early perceptual grouping which sets up the candidates for object identification. It implies that preattentive organization exists only within dimensions — within a color map, within a shape map, within a map of movements or orientations — and that these maps are related to each other only where and when attention is focused. This suggests the possibility that we might effectively camouflage an object at this preattentive level by placing it at a boundary between two groups, each of which shares one of its features. We can choose an object which within either group alone would be quite salient, and see if adding the second group makes it harder to see. Figure 5 shows an example. In order to detect the presence of the red 'X', attention has to be narrowed down to exclude the adjacent red 'Os' and blue 'Xs' and focused on the item itself. Yet in either group alone, it would be quite easy to detect. If, on the other hand, the target has a unique feature (for example the color green or vertical lines), we would expect detection to be independent of grouping. We ran an experiment to test these predictions, and the predicted camouflage of the conjunction target was clearly confirmed: subjects took 135 msec longer to find the conjunction target than the feature targets. In fact, they missed it altogether on 9% of trials, even though the display remained on until they responded. What seems to happen is that two competing ways of grouping this display exist — one within the preattentive color map and one within the preattentive shape map. The conjunction target exists in neither of these maps, while the feature target is always unique in one of the two. Even when we place the targets in the center of a group instead of at the boundary, it still takes longer to detect the conjunction than to detect the feature target. The presence of distractors elsewhere in the display which share the locally distinctive feature of the conjunction target forces us to narrow attention down, at least to exclude these irrelevant distractors. Luckily in normal life the preattentive boundaries of our multiple feature worlds are likely to agree. The features of real physical objects have highly correlated spatial boundaries. The edges of the dog co-exist and move together, whether we define them by their color — brown, or by their texture — furry.

What is the next prediction? I have been claiming that we cannot identify a conjunction of target features without locating it in order to focus attention on it. This need not be true of simple features. For these we may be able to detect their presence or identify them without first locating them. In fact, the occurrence of illusory conjunctions suggests that features could be to some extent free-floating, not precisely pinned down or labeled with their spatial locations. In order to see boundaries between areas with different features, we need to locate the position of the discontinuity. But within any perceptual group we may not know, at the feature level, where any particular item is. Of course we can rapidly find its location when we have detected its presence. But my hypothesis was that this might require an extra operation, which might take a measurable time.

I set up an experiment to test this idea, simply by looking at the dependencies between the two types of judgment, identification and localization. If we are unable to identify conjunctions, without locating them, the dependency ought to be complete. In other words, we should never identify a conjunction without also knowing where it is. This should not be the case for features, if locating them is actually a separate operation from identifying them. The task required subjects to make a forced choice identification of which of two targets was presented in an array of 12 colored letters. In the conjunction condition, the targets were a red 'O' or a blue 'N' and the distractors were red 'Ns' and blue 'Os'. In the feature condition, the targets were an 'S' or an orange letter with the same distractors as in the conjunction condition. We set the exposure duration (followed by a mask) to get about 80% correct identification of the target in each condition. We then asked subjects not only to identify but also to locate the targets, by writing their answer in a matching grid. We found that subjects were very unlikely to identify the conjunctions without also knowing where they were, while this happened on about 40% of trials with feature targets. We got the same result when we matched the exposure durations instead of matching the error rates. Again, we have evidence that detection of a feature can occur not only without focused attention, but also without information about its spatial location.

At some level we may have "free-floating" presence or absence information for features but not for conjunctions.

Let me summarize the conclusions so far. I think we have found some support for most of the predictions I made earlier. I would like to emphasize the importance of their convergence rather than putting too much weight on any one alone. It seems that we can identify separable features in parallel across a display. We do this separately within a number of independent, pre-attentive feature maps. Relating any individual feature to other features of the same object requires an additional operation. This feature registration can mediate texture segregation or figure ground grouping, and can locate potential objects to be identified by a serial scan with focal attention at the next level up. Identifying conjunctions of features on the other hand, requires focal attention and serial scanning of locations; it therefore cannot mediate texture discrimination. This finding is largely independent of spatial density, difficulty and practice, although these variables may change the rate of serial scanning. The conclusions seem to apply not only to feature values on different dimensions (color and shape) but also to local elements of shapes, lines, curves, and angles). Thus we have converging evidence from a number of different paradigms and stimuli all of which so far supports the theory.

However, you may have noticed a rather glaring omission. I didn't venture to test the central claim of the theory — that illusory conjunctions would occur if we prevented focused attention. The reason I hesitated, of course, was that I didn't really believe they would occur. Nevertheless, shame eventually began to prevail over my fears, and I decided I could no longer put off doing the most obvious experiment. So we devised a pilot experiment to try at least on ourselves in the privacy of the testing room. We used displays like those in Figure 6. The primary task was to attend to and report the digits; we then attempted to write down all we had seen of the colored letters. We stopped, discouraged after a few trials, convinced we were seeing correctly the one or two items that we managed to report. Each of us found it hard to believe she had written down almost as many illusory conjunctions as correct items. Having clearly seen a pink 'T', we were reluctant to accept the evidence on the card which contained a pink 'X' and a green 'T'. I tell this anecdote because it is important to distinguish whether I'm describing just wrong guesses, memory failures, or whether these responses reflect genuine perceptual experiences. We certainly believed we had "seen" at least some of the illusory conjunctions rather than guessed them in the absence of information. I would like to claim that what we are studying reflects the way we construct our mental experiences of the outside world.

We tested other subjects on the same displays, to see whether they replicated our results, asking them to report only what they saw or were reasonably confident they saw — not to guess. We found that most did make a large number of conjunction errors, in fact almost as many as correct reports. They averaged about one illusory conjunction in every two trials. If we compare these errors with those in which they got one feature correct and the other a feature intrusion which had not appeared on the card, we find they were three times as likely. Again, it is worth mentioning another anecdotal observation, in relation to the question whether subjects were guessing or 'seeing' the illusory conjunctions: several subjects stopped after a few trials and spontaneously made comments like 'Oh, you are tricking me. The numbers were colored that time'. No "demand characteristics" or response bias will explain that unrequested observation, since we had told subjects that all the digits were black.

The next question we asked was whether there are any constraints on these illusory recouplings. For example, can we take the red from a small, outline circle and use it to fill in the area of a larger triangle that was originally blue? Or are there limits to the sizes, shapes, areas and distances between which the exchanges can occur? The letters in the first experiment were the same size; did this encourage the promiscuous mixing of colors and shapes, or would we get it as freely with more heterogeneous displays? We tried displays in which we deliberately varied the color, size, shape and whether the color was filled in or outlined (as in Figure 7). Some displays varied in only two features, some in three and some in all four.

We also ran two different attention conditions. It is important to the theory that illusory conjunctions should result from attention failures, and not from any other form of difficulty. We therefore compared this divided attention with a focused attention condition. In the divided attention condition, subjects attended to and reported the digits as their primary task, and were cued only after the display was presented which colored shape to report. In the focused attention condition, on the other hand, the digits could be ignored, and subjects were given a spatial cue 150 msec in advance of the display, telling them which colored shape to attend to and report. We matched the overall accuracy in the two conditions by reducing the exposure duration in the focused attention condition until the stimuli were hard to see. To summarize the results: (1) first in the divided attention condition, all four features were liable to switch. Subjects reported illusory exchanges of size and outline versus filled as well as of shape and color. All four dimensions seem to be separable by this criterion. (2) Illusory conjunctions were as frequent with heterogeneous displays as with more homogeneous ones. Colors were as likely to switch between items which differed on the other three features as between items which were otherwise identical. (3) Illusory conjunctions do seem to be linked to attention load rather than to discriminability or task difficulty in general. When we changed from post-cued, divided attention to pre-cued, focused attention, we found a big shift in the type of errors. With attention cued in advance and no primary load, subjects made very few conjunction errors; either they omitted the item altogether, or they were as likely to report a feature that was not presented as a feature from the wrong location in the display.

These last three experiments have confirmed the initial premise of the theory. To a first approximation, each feature seems to be coded as an independent entity; it can migrate without constraints from its source or destination, if attention is diverted elsewhere. An object can be as confidently seen when its conscious representation is generated from the color and size of one object and the shape and filled-in property of another as when it accurately matches the features of a physically presented stimulus. Features are exchanged as freely between objects which differ maximally as between otherwise identical objects, even though this usually requires a change in the conscious representation of the migrating feature. For example, moving a color between objects of different sizes or between one outline and one filled object must also change the amount of color perceived.

The implications of this conclusion, if we accept it, are quite far-reaching. It suggests that the internal representation, on which conscious experience depends, contains discrete labels of values on each dimension separately rather than a wholistic, interactive record. The whole object must be resynthesized from a set of these discrete feature labels, which may have been accidentally interchanged. If in a brief glance only the labels 'blue', 'small' and 'triangle' are registered, we supply our conscious image with the correct quantity of blue coloring to fill the specified area, regardless of how much was initially presented. This hypothesis involves an extreme interpretation of the notion of feature separability, and places conscious seeing at a greater remove from the physical stimulus than we would intuitively assume. I believe the Gestalt psychologists were right in claiming that consciousness is peopled not by disembodied features but by objects, their backgrounds and their interactions. But in terms of the operations which *mediate* this perception, the elements which code the necessary information appear to be abstracted from their contexts; discrete and independent labels which require attention for their correct resynthesis.

The last question I want briefly to touch on is one that I'm sure has occurred to many of you as you listened. So far, I have talked as if we knew quite clearly which properties count as separable features and which have to be put together as conjunctions. In fact, of course, this is far from obvious, and is, I believe, an empirical question. My claim has been that attention is needed whenever we have no population of specialized detectors which directly sense a particular feature in the display. Does this put me into a vicious circle in which I define a separable feature as one which requires attention before it can be correctly integrated with others? One answer is yes, but it is not the whole answer. The escape from the circle is through a variety of converging tests for separability, all of which, we hope, will pinpoint the same candidates for separable feature-hood. The strategy is to choose two features which are most likely to be separable — for example we might use color and line orientation — to establish across a variety of tasks two different behavioral syndromes typical of features on the one hand and conjunctions on the other. So, just as a physician uses spots, a fever and a sore throat to diagnose measles, we can then use the feature syndrome — parallel search, location errors, illusory conjunctions, texture segregation — as new diagnostic criteria for separability to add to those proposed by Garner. We may then be able to apply these tests to help us decide, for example, whether more dubious features like closure or intersection or symmetry qualify as elementary features, whether faces are recognized as unitary wholes or built out of features, and so on. We can also look at perceptual learning to see whether new detectors can be set up, registering conjunctions of previously separable features as unitary wholes.

With color and shape, there may be built in constraints on unitization. For component parts of shapes (e.g., curves, lines, angles, etc.) it would make biological sense to allow some flexibility. We want to be able to develop automatic detection of important faces and places. Another difference between color-shape conjunctions and shape component conjunctions is that the latter may produce emergent features (Pomerantz et al., 1977). For example, conjoining L with \ produces a triangle Δ, which has a new property, closure. Are these emergent properties also picked up by separate populations of feature detectors? If so, they could affect the probability of illusory conjunctions. People might be less willing to see a Δ given an incorrect conjunction of L and \, simply because the closure feature was missing. The more salient an emergent feature is, the more its absence should inhibit the formation of an illusory conjunction. One test of this question is to compare two conjunctions of the same features, the diagonal line and the right angle —one which generates closure and one which does not (Δ versus ∠). We are finding more illusory conjunctions with the arrow than with the triangle (24% vs. 15%). Closure may be more salient as an emergent feature than intersection. But this generates another prediction: if there is an emergent feature of closure which is picked up by a separate population of detectors, it should also behave like a feature in the other tests: it should allow parallel search and mediate texture segregation. So search should more often be serial with arrows than with triangles (see Figure 8). This prediction is confirmed. The slope of the search function was steeper and more linear with arrows, and the mean ratio of positive to negative slopes was close to 0.5 for arrows (0.47) but not for triangles (0.22). Texture segregation was also slower with arrows than with triangles, averaging 990 msec versus 786 msec. So, across stimuli, the tasks seem to covary as predicted.

Another prediction is also possible: There were quite marked individual differences, particularly with the arrows. This suggests the possibility that different individuals may give different weight to different features in determining whether to 'see' an arrow or a triangle. For some people, three lines in the right arrangement might be enough, while for others, closure may also be essential. This gives us another way of testing the theory; of seeing whether the different tasks I've discussed are related in the way I suggest. Just as different stimuli may generate performance patterns which covary across the tasks, so may individual subjects with any given stimulus. People who make illusory arrows out of their parts should also search serially for arrow targets in conjunction distractors, and should have trouble segregating textures which contain arrows or their parts. On the other hand, people who code arrows by some emergent wholistic property, for example the three-way intersection, should find an arrow target in parallel across search displays and should be able to use the presence or absence of arrows to segregate one area from another.

We tested these correlations across tasks within each type of stimulus. The findings were that for the arrows, there were significant correlations within individuals between (a) the frequency of illusory conjunctions, (b) the difference in search slopes for arrow targets in a conjunction and in a feature background, and (c) the difficulty of texture segregation. For triangles, the correlations were not significant, perhaps because so few illusory conjunctions were seen. It looks, then, as if closure does function as a primitive feature in the same way as line orientation and color. It also looks as if the three tasks I used are inter-related in the way the theory suggests. This is certainly an example of boot-strapping Psychology, but, I would argue, not as such inadmissable.

Finally, is there a personality trait of proneness to illusory conjunctions? If you make arrows from their parts, do you also tend to make triangles from theirs? The answer is 'yes': The within-subject correlation of 0.5 was significant. Whether it would generalize to color-shape conjunctions is another question.

That brings me to the end of the data I want to describe today. I will return for a last few words to the problems this account may raise for our everyday perception of objects and complex scenes, or of words, sentences and meanings in reading. Can we reconcile what I've been claiming with the apparent speed and richness of information processing that we constantly experience? I can only speculate. Perhaps this richness at the level of objects and scenes is largely something we create, an informed hallucination. We can certainly register a large array of features in parallel, and we can do this along a number of dimensions. But, if we apply more stringent tests to see how accurate and detailed we are in putting them together, e.g., in perceiving faces or words, the results tend to be less impressive. Scanning a school photograph for my daughter's face, among hundreds of other teenagers is a painstaking business. Proof-reading is also best done slowly and serially. In both these examples, contextual redundancy is less useful than it normally is. Much of our peripheral or non-attentive 'seeing' may capitalize on our prior knowledge. It may consist of matching expected features to actual features without checking on how they are combined. If so, the 'wholes' or objects would still exist in our heads. They would be what we expect to see and normally end up seeing. Yet they may not be the initial code that registers the stimulus in its sensory form. I suggest there may be three ways in which we can see whole objects, and we may not be aware which we have used in any given instance. 1) We may see them, as I've suggested, by integrating their features in the spotlight of attention. 2) We may see them by predicting their features in a familiar context and separately confirming that each feature is present. 3) Finally, in the absence of either prior information or focused attention, we may be reduced to random resyntheses which result in illusory conjunctions.

References

Gardner, G. T. Evidence for independent parallel channels in tachistoscopic perception. *Cognitive Psychology*, 1973, *4*, 130-155.

Garner, W. R. *The Processing of Information and Structure*. Potomac, MD: Lawrence Erlbaum, 1974.

LaBerge, D. Attention and the measurement of perceptual learning. *Memory and Cognition*, 1973, *1*, 268-276.

Pomerantz, J. R., Sager, L. C. & Stoever, R. J. Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance*, 1977, *3*, 422-435.

Shiffrin, R. M. & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending and general theory. *Psychological Review*, 1977, *84*, 127-190.

Zeki, S. M. The functional organization of projections from striate to prestriate visual cortex in the rhesus monkey. *Cold Spring Harbor Symposia on Quantitative Biology*, 1976, *15*, 591-600.

SEARCH FOR COLORED SHAPES

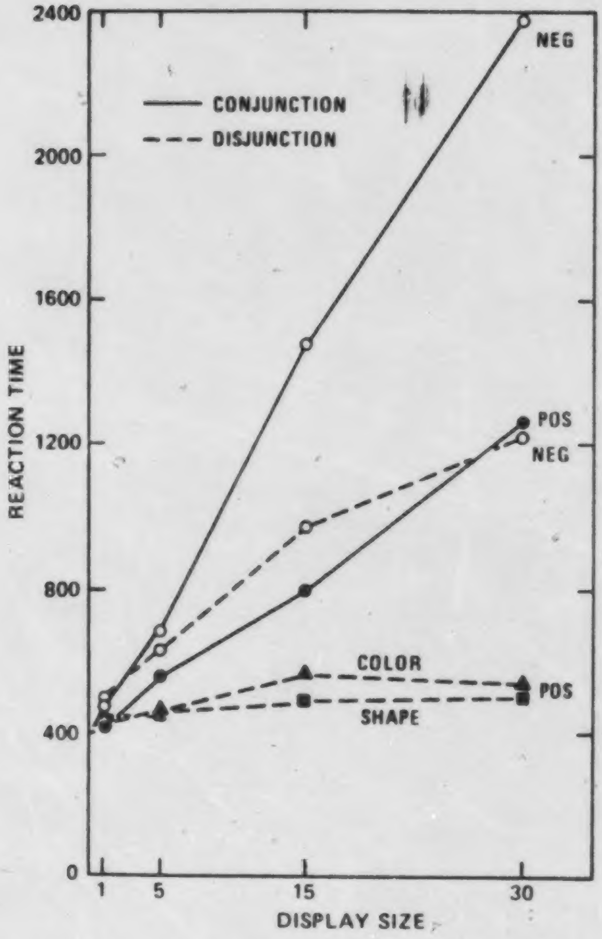

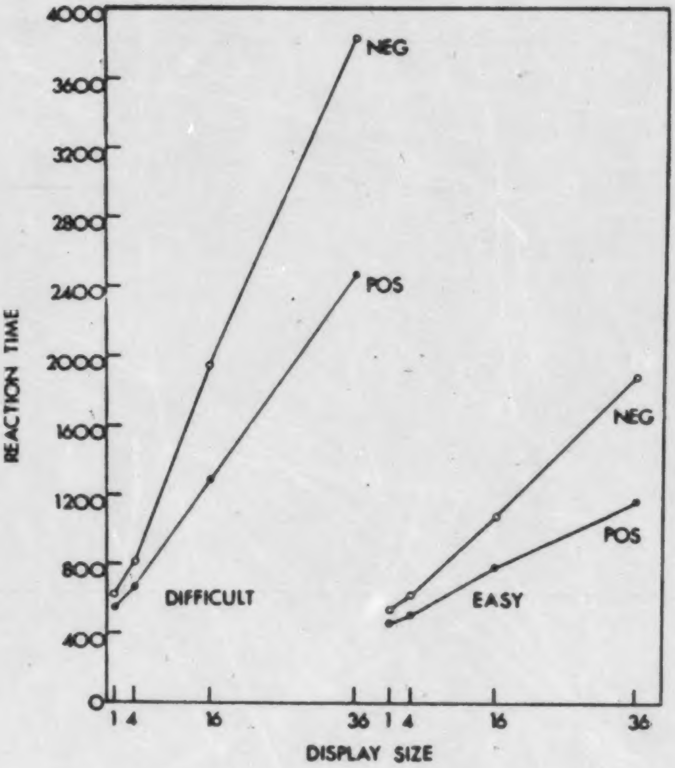
4000
3600
3200
2800
2400
2000
1600
1200
800
400
0
REACTION TIME
NEG
POS
DIFFICULT
NEG
POS
EASY
1
4
16
36
1
4
16
36
DISPLAY SIZE

SPATIAL DENSITY AND SEARCH FOR CONJUNCTIONS

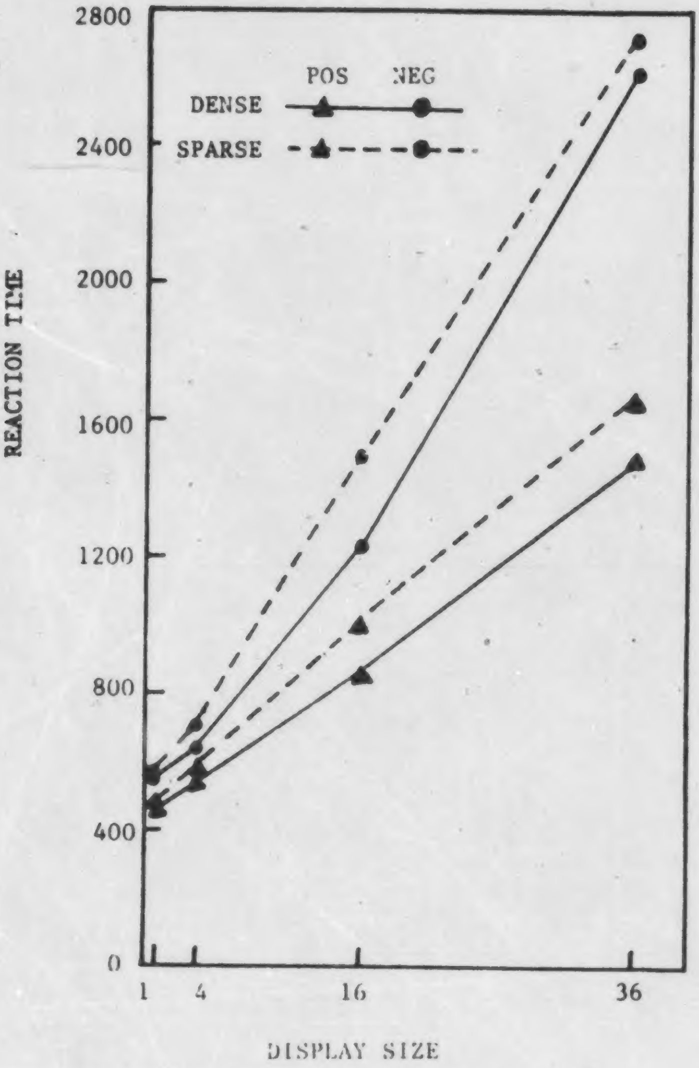

255

LETTER SEARCH

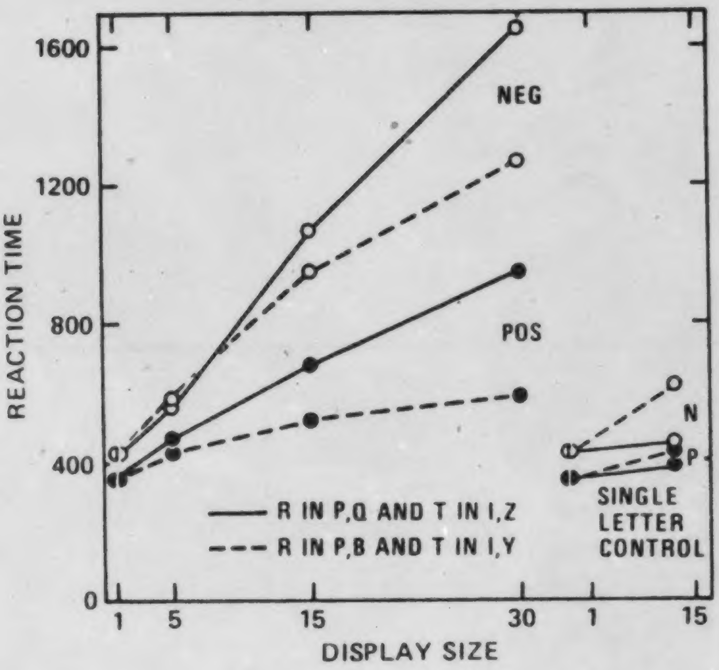

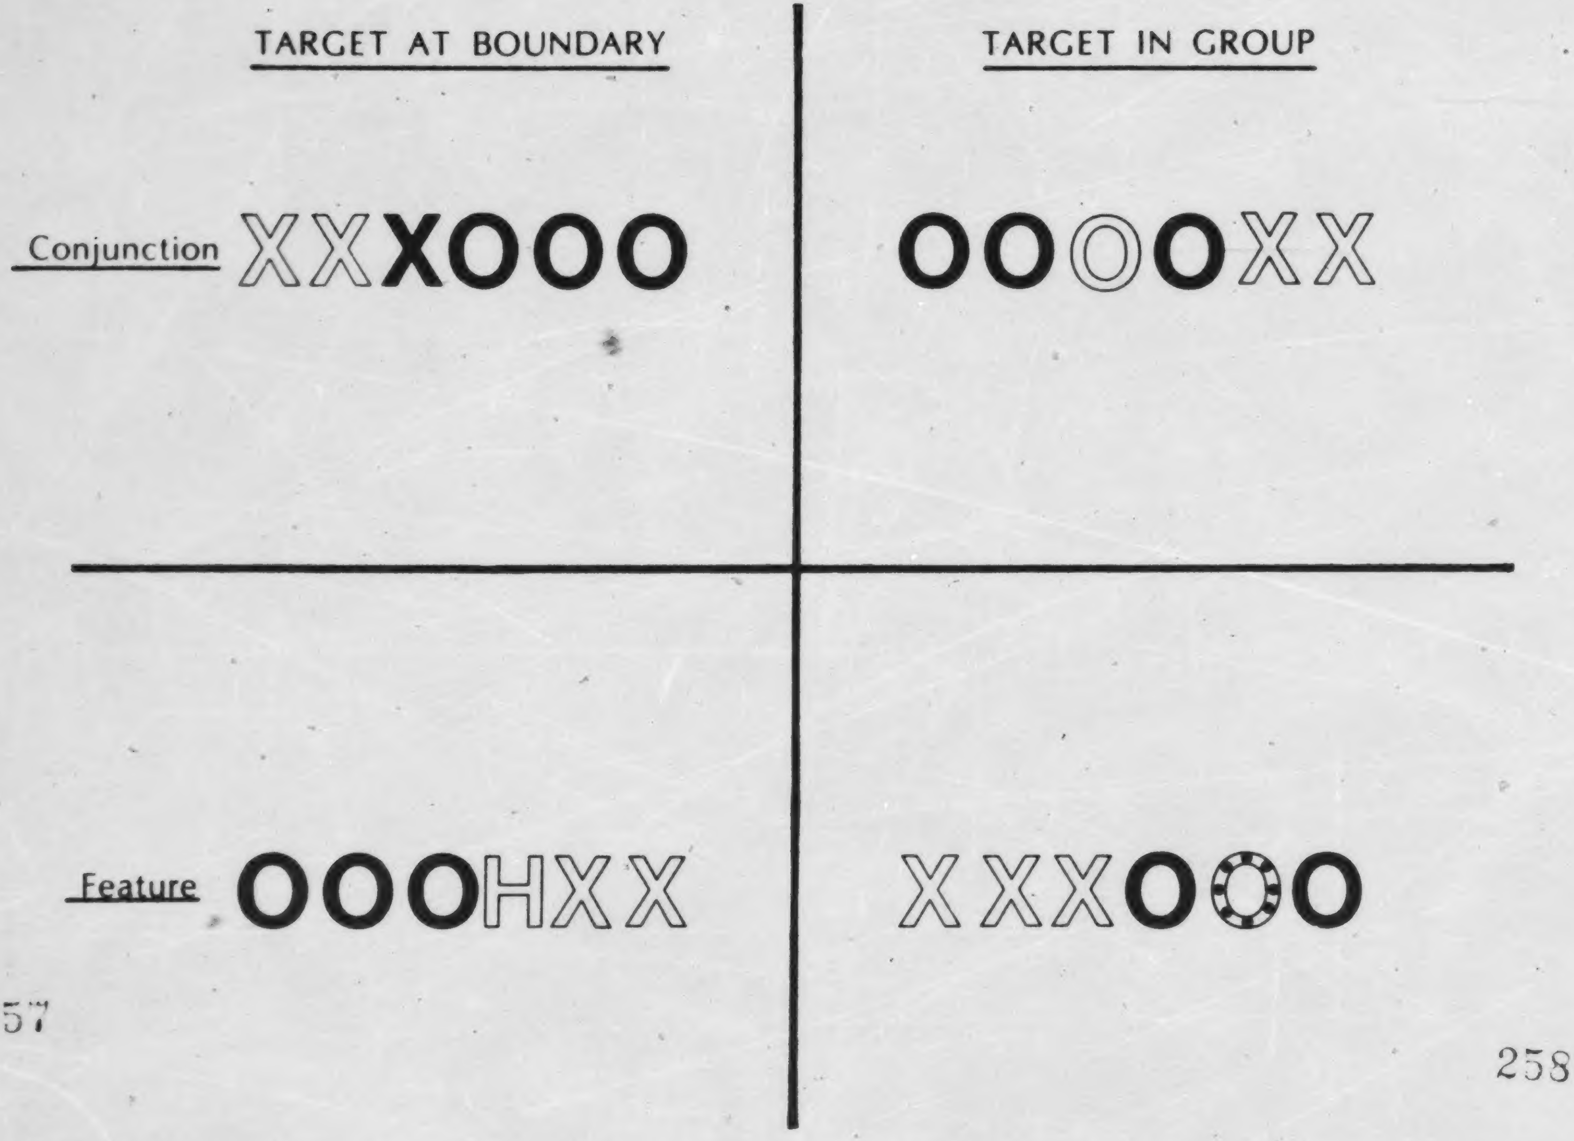
TARGET AT BOUNDARY
TARGET IN GROUP
Conjunction
XXXOOO
OOOOXX
Feature
OOOHXX
XXXOOO

DISPLAY USED TO DEMONSTRATE ILLUSORY CONJUNCTIONS

4 N T O 7

The digits were black and each letter was a different color.

6

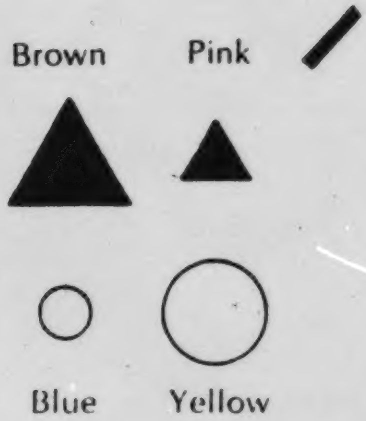
Brown
Pink
Blue
Yellow

7

EXAMPLES OF DISPLAYS USING TRIANGLES AND ARROWS

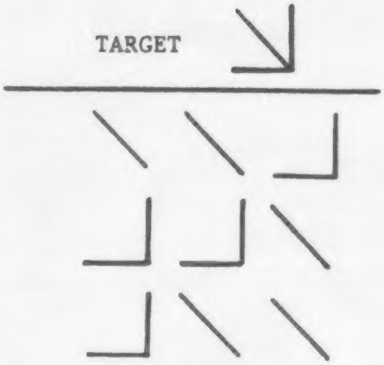

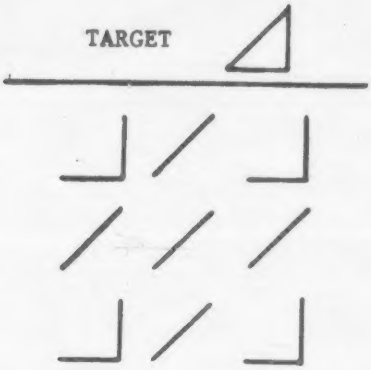

Consequences of Visual Orienting[1]

Michael I. Posner and Yoav Cohen
University of Oregon

I. Introduction

During the past several years we have been studying how peripheral visual events produce covert orienting of attention. Much of this work has now been published (Posner, 1980; Posner, Davidson & Snyder, 1980; Remington, 1980; Shulman, Remington & McLean, 1979). These studies show covert attention movements are sufficiently time-locked that we can trace their course in terms of changes in the efficiency of responding to probe stimulus events that occur at different places in the visual field.

In this paper we ask: What are the consequences of having oriented to a peripheral cue once attention is returned to fixation? We have two basic reasons for asking this question. First is an intrinsic interest in the dynamics of what influences the likelihood of attending to a spatial position. Do we tend to interrogate a previously active area again? Or do we tend to avoid it in favor of a fresh source of information? Repeated events are known to habituate, but repetition can also yield faster, more efficient response times. Second, work on orienting to letters by McLean & Shulman, 1978, has shown facilitation of the attended pathway--even after attention is withdrawn. Presumably, attending to a letter enhances the efficiency with which it can be activated by input even after attention is reoriented toward another task. We wish to compare the results obtained with orienting to visual location with the findings based on orienting to higher level codes.

---

[1] Paper delivered to Conference on Attention, Seattle, September 1980 and to the Psychonomics Society, November 1980. This research was supported by the National Science Foundation Grant #BNS 7923527.

II. Basic Paradigm

Our subjects look at a cathode ray tube display consisting of a central box which is flanked by two peripheral boxes placed 8° to the left and right of the fixation box. The trials begin by 150 millisec brightening of one of the two peripheral boxes selected at random. A bright probe dot occurs either 0, 100, 200, 300 or 500 milliseconds following the brightening. The dot is usually in the center box (.6), but it may occur either side (.2). Catch trials in which no probe is presented occur with probability of .2. Subjects are instructed to respond to the dot as quickly as possible by pressing their single key. Eye movements are monitored using EOG electrodes and trials with detectable movements are excluded. The first study used six Ss run on two days.

Our expectation was that the cued stimulus would summon attention. Thus, the cued side would have an initial reaction time advantage over the uncued side. However, because probes occur mainly at the center, subjects should try to keep attention there insofar as they can. Thus, the initial advantage to the cued side should be lost as full attention is given to the center. We can then compare the two sides to observe the consequences of the previous facilitation by the cuing.

The results conform well to our expectation. These are shown in Figure 1. The cued side shows an advantage for the first 150 milliseconds. This is replaced by a clear inhibition after about 300 milliseconds. The center remains fastest throughout as might be expected both by its high probability and its foveal location. It looks as though the consequence of the early advantage to the cued side is a subsequent inhibition.

Insert Fig. 1 About Here

263

Several views about the origin of this inhibition occurred to us: 1) It could be due to the fact that this is a comparison of two alternative positions. In many two-alternative reaction time tasks alternations turn out to be better than repetitions. Failing to find a probe at the cued position shortly after the cue perhaps the subject guesses that the probe is more likely to occur at the other position. 2) It could be due to the movement of attention away from the cued stimulus in order to return to the center. If the subject moves attention back toward the center, perhaps he has more difficulty in reversing it back to the cued position than he does in allowing it to continue to the uncued side. 3) Some part of the pathway from the cued location is reduced in efficiency by the cuing. This could occur because of the sensory cue itself, or because of the orienting that occurs as a result of the sensory cue.

III. Four Alternative Experiment

In order to test the first two ideas, we used a central box and four peripheral boxes, each 5° from the central box. The probe occurred at the central box with probability .6 and in one of the peripheral boxes with equal probability of .1. Otherwise, the experiment was the same as the previous one.

Figure 2 shows the results for 12 subjects. Comparing the cued position with the mean of the other three positions there is an initial advantage for the cued side replaced by an inhibition as before. When the side opposite the cue (far position) is compared with the two orthogonal (near) positions, it is clear that the far position is no faster than the two near positions. These results show that the inhibition is not limited to the two alternative case

----

Insert Fig. 2 About Here

----

and that a stimulus position in the direction of the assumed attention movements from cue to center is not necessarily at an advantage over other positions in the visual field. The results thus eliminate the first two explanations for the inhibition effect.

IV. Double Cuing

In order to examine the role of sensory factors in this phenomenon we introduced cuing either by brightening the cued box or by dimming it. If the facilitation effect is indeed attentional and not due to forward brightness enhancement, for example, we should get similar facilitation in both cases. To check further on the sensory versus attentional character of facilitation and inhibition we introduced trials in which both peripheral positions were cued simultaneously. To summon attention back to the center we brightened the center position 300 milliseconds following the initial cue. No probes were presented at the center but they occurred with equal probability at the two flanking boxes, either 80 milliseconds following the cue or 500 milliseconds following the cue. In accordance with the previous result we expected the cue side to show facilitation in the former condition and inhibition in the latter.

The results are shown in Figure 3. Data from the single cue trials conform well to our expectations. Regardless of whether the cue was introduced by brightening or dimming the cued side is initially faster than the uncued and is slower for probes 500 milliseconds following the cuing. This suggests the facilitation effect is not due to any kind of brightness enhancement.

The results of the two cue condition are more interesting. The cued sides are not significantly facilitated when compared to the uncued side in single cue trials. This suggests, in accordance with previous work, that attention cannot be split to the two sides when both are cued. However, the inhibition

Insert Fig. 3 About Here

in the two cue trials is as great as that found for single cue trials. This suggests that the inhibition effect does not arise from attentional orienting but from the information presented at the cued position.

V. Arrow Experiment

To check further on the role of sensory information and attentional orienting in facilitation and inhibition we used a central cue rather than a peripheral cue to indicate where to attend. Each trial began with an arrow that occurred above the central fixation box. The arrow indicated that the cued side would have a probe with probability .8, while the uncued side (opposite the arrowhead) would have probes with probability .2. 600 milliseconds after the arrow, attention was returned to the center by brightening the center box. Probes following the center brightening are most likely to occur at the center (.6) and have probability .2 of occurring on either side. Probe events occur either at 450 milliseconds after the arrow when attention should be on the side cued by the arrow, or at 950 or 1250 milliseconds following the arrow when attention should have been returned to the center by our logic.

The results are shown in Figure 4. The cued side shows the expected

Insert Fig. 4 About Here

facilitation following the arrow cue. This is in accord with many other results obtained on central cues. However, there is clearly no inhibition following the return of attention to the center. Thus the inhibition effect, but not the facilitation effect depends upon presenting sensory information at the periphery.

VI. Conclusion

Our most direct conclusion is that visual information both summons attention and serves to inhibit the processing of further information at that place in space. Under many circumstances the presence of attention will compensate for the habituation caused by the cue, but attention seems to have to work against a basic sensory bias that favors fresh information sources. This complex but exquisite reciprocal relationship between sensory information and attention presumably arises in the need to be able to sustain information to a source of repeated signals and the need to be able to summon resources to fresh signal sources. It warns us that although any given experiment may be aimed at the study of sensory or attentional processes, careful control is needed to tease apart the contribution from each level. Many experiments, supposed to be purely sensory, must involve a net effect of both sensory and attentional processes. Now that we have methods of the control and time-locking of attention to aspects of the signal, it may be possible to understand more fully the contribution of these two levels to the overall information processing.

The results obtained with orienting to visual positions in space seem at first quite different from those that McLean & Shulman (1978) reported for attention to letters. When attention is given to a letter it remains active and thus is more efficiently processed than any other letter once attention is withdrawn. We have never observed inhibition in the letter experiment. Perhaps at a deeper level similar reciprocal relations between purely sensory activation of a semantic pathway and attention will appear in further experiments. We have not yet been able to locate a model for the thorough examination of this possibility.

We have begun to apply insights from the reciprocal relation between attention and sensory processes to two areas. First, we are trying to understand how people are able to maintain concentration on a source of signals for an extended time. Put in its usual terms: How does filtering occur? Our results suggest that filters require an active orienting of attention to overcome the advantages that accrue to fresh sensory pathways. Second, we are tracing the effects of midbrain and cortical injury on these components of orienting. The ability to time-lock the operation of cognitive mechanisms of attention, as shown in our peripheral and central cuing studies, provides a rich methodology for exploring attention, even in patients who have difficulty understanding more complex instructions. We expect these methods to be useful in teasing apart the underlying neural systems that subserve aspects of visual orienting.

268

References

McLean, J. P. and Shulman, G. L. On the construction and maintenance of expectancies. Quarterly Journal of Experimental Psychology, 1978, 30, 441-454.

Posner, M. I. Orienting of attention. Quarterly Journal of Experimental Psychology, 1980, 32, 3-25.

Posner, M. I., Snyder, R. R. and Davidson, B. J. Attention and the detection of signals. Journal of Experimental Psychology: General, 1980, 109, 160-174.

Remington, R. Attention and saccadic eye movements. Journal of Experimental Psychology: Human Perception and Performance, 1980, 6, 726-744.

Figure Captions

Figure 1 Reaction time to a probe dot as a function of cue to probe SOA for probes on the cued and uncued side and center.

Figure 2 Reaction time to a probe dot as a function of cue to probe SOA for probes at the cued position, position opposite the cue (far) and two orthogonal (near) positions and the center.

Figure 3 Reaction time to a probe dot following a cue consisting of brightening (left panel) or dimming (right panel) the surrounding as a function of cue to probe SOA (interval) for cued and uncued sides and trials on which both sides are cued (both).

Figure 4 Reaction time to a probe dot as a function of SOA (interval) following a central arrow cue for cued, uncued and center positions.

269

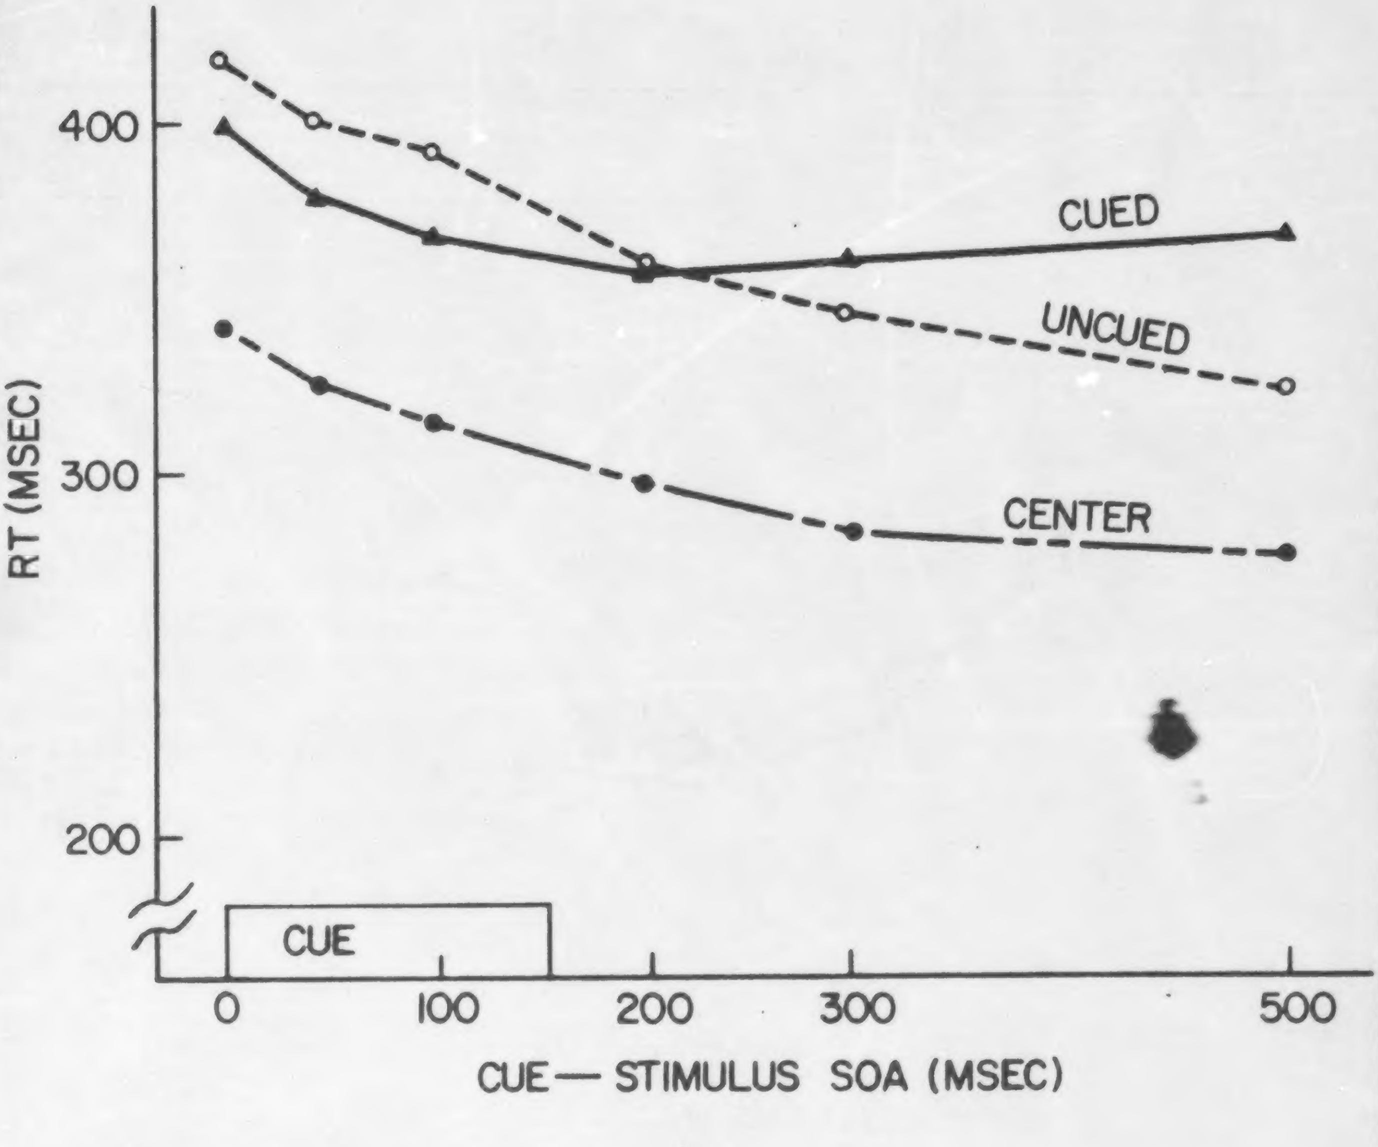
CUED
UNCUED
CENTER
RT (MSEC)
400
300
200
CUE
0
100
200
300
500
CUE— STIMULUS SOA (MSEC)

270

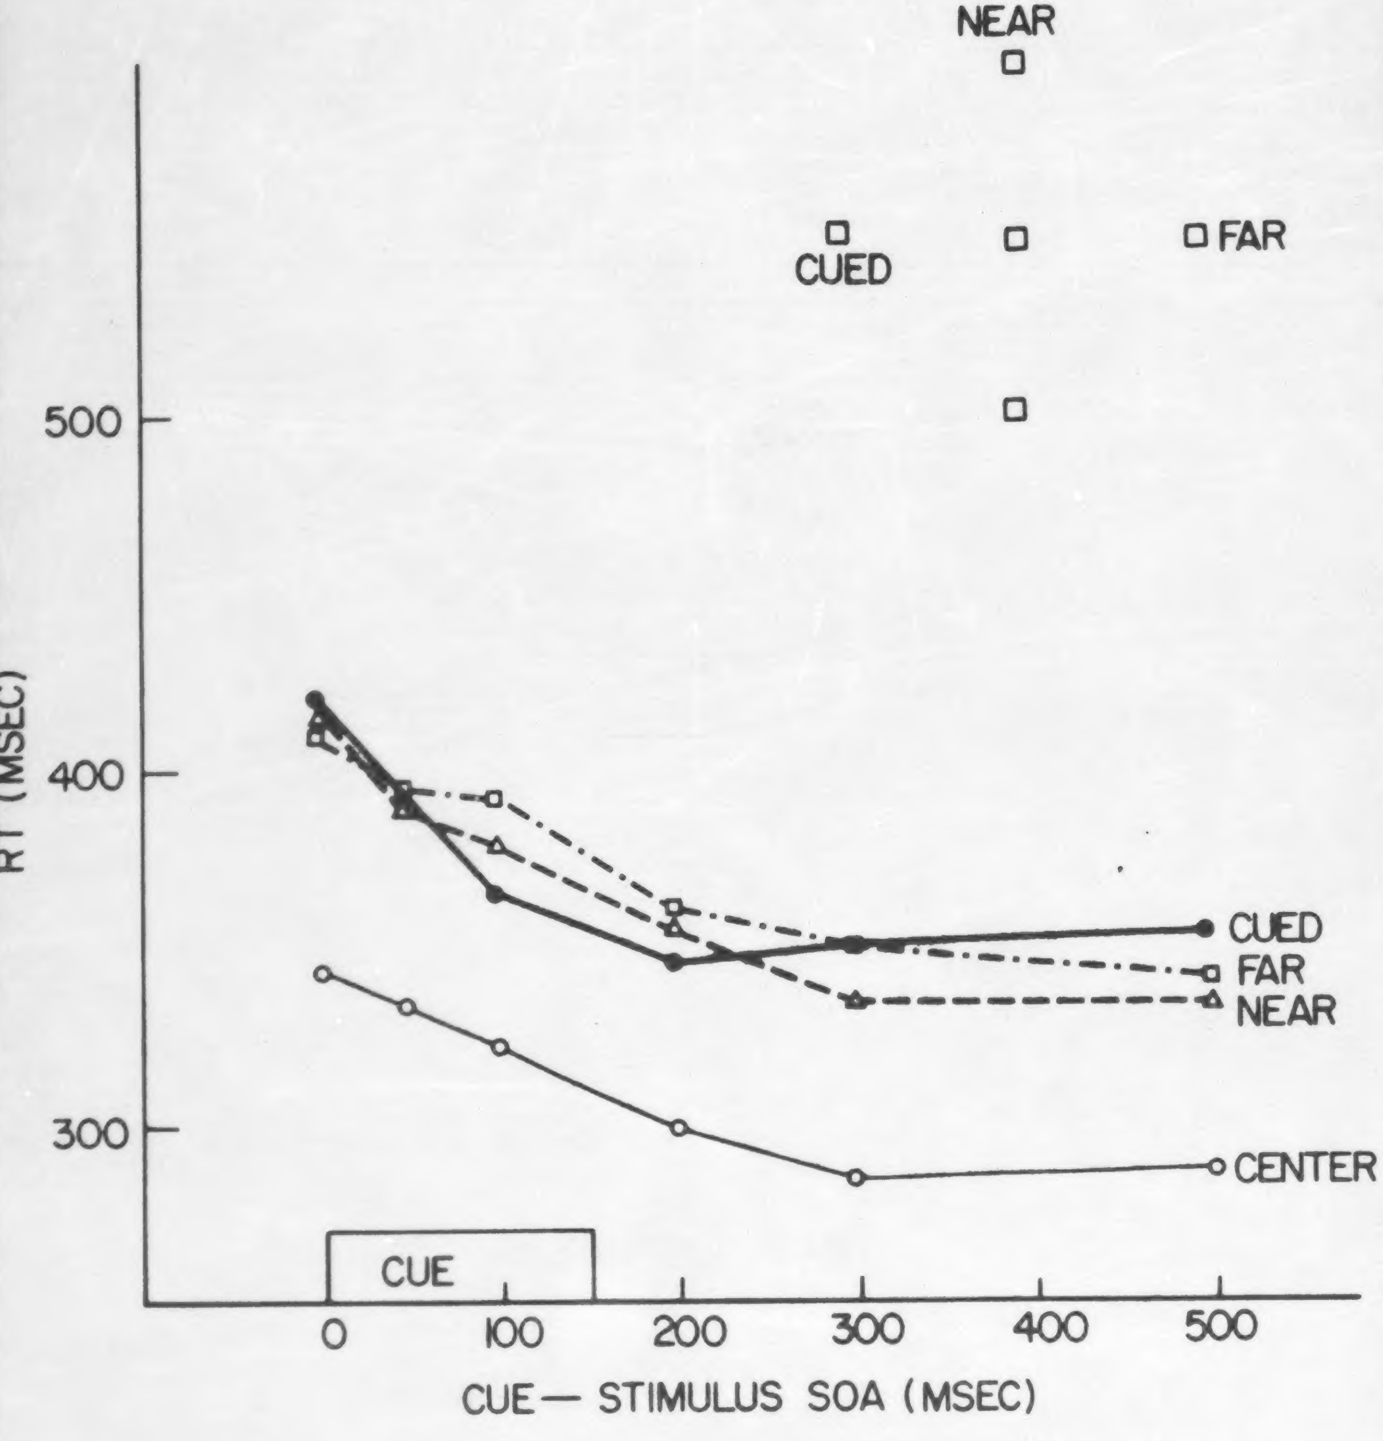
NEAR
CUED
FAR
500
400
300
RT (MSEC)
CUED
FAR
NEAR
CENTER
CUE
0
100
200
300
400
500
CUE— STIMULUS SOA (MSEC)

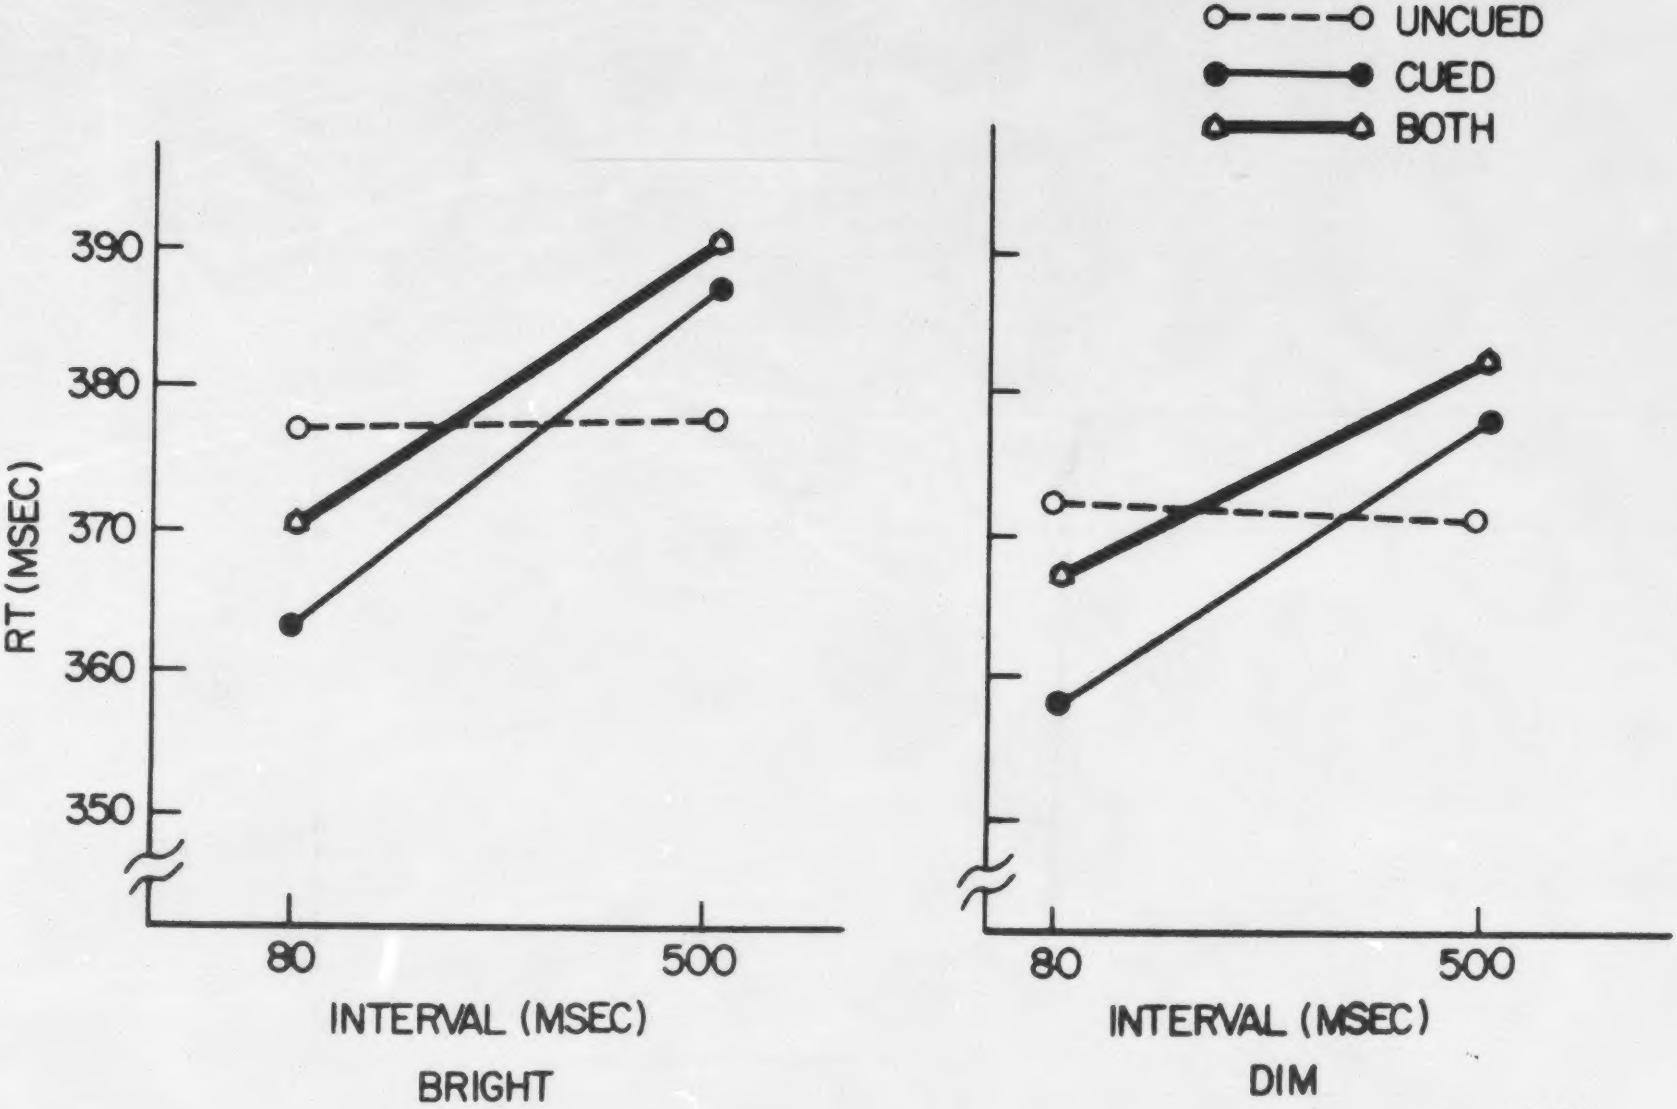
UNCUED
CUED
BOTH
RT (MSEC)
390
380
370
360
350
80
500
INTERVAL (MSEC)
BRIGHT
80
500
INTERVAL (MSEC)
DIM

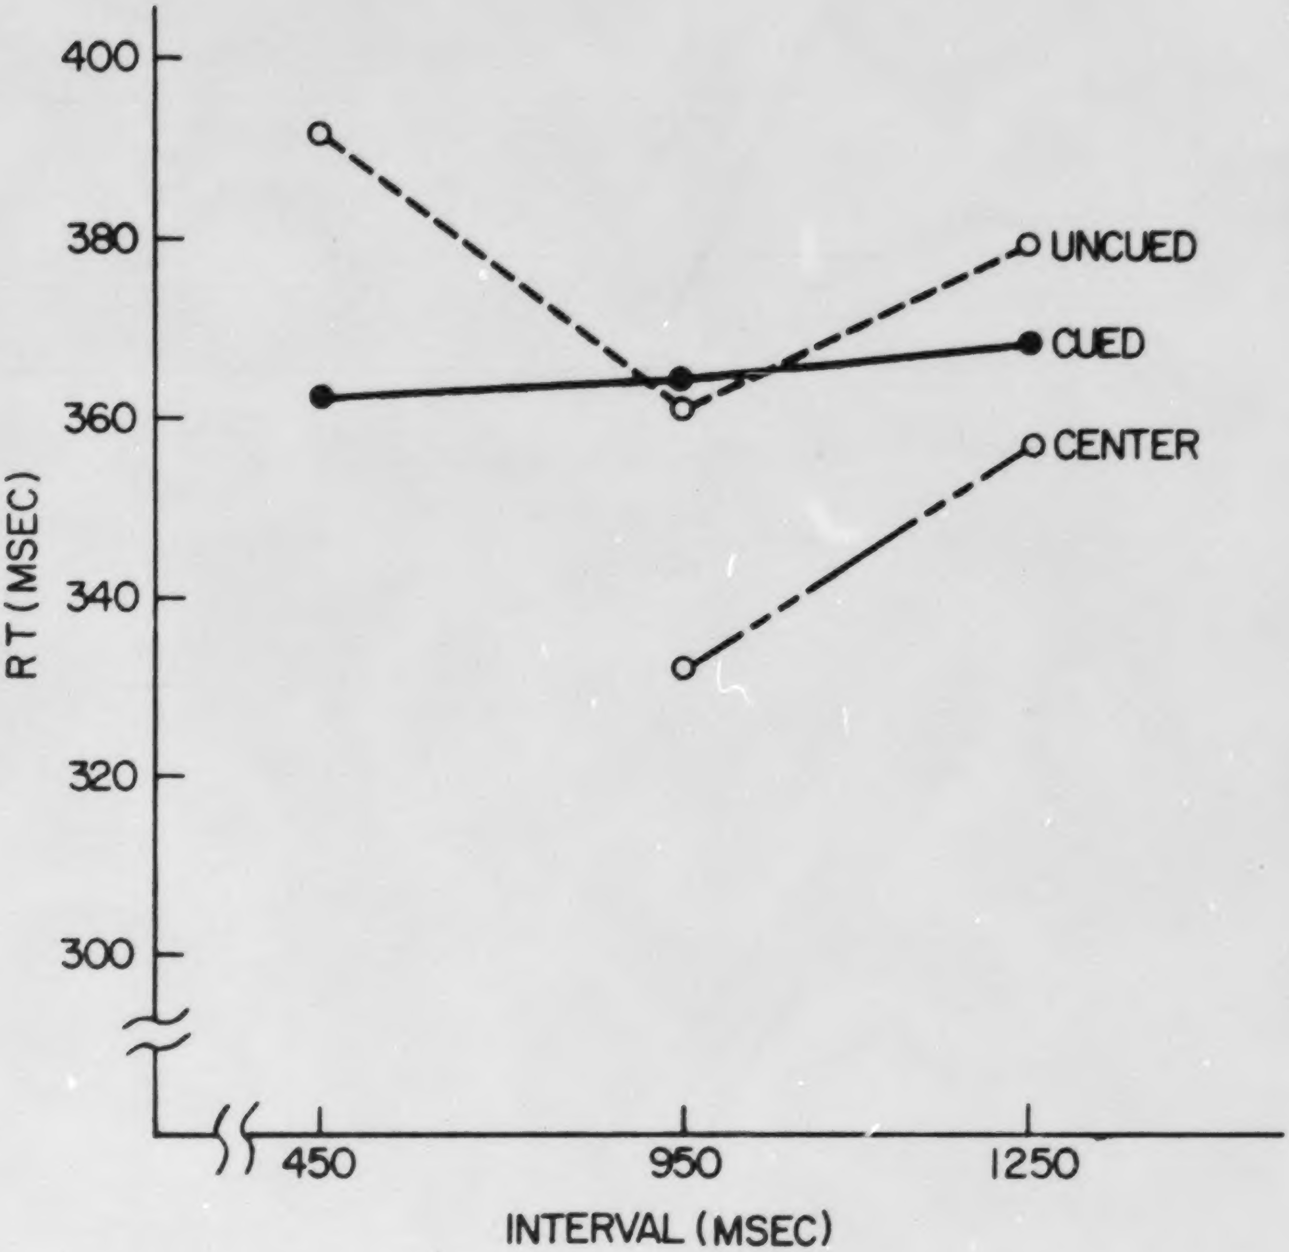
400
380
360
340
320
300
RT (MSEC)
450
950
1250
INTERVAL (MSEC)
UNCUED
CUED
CENTER

UNITIZATION AND AUTOMATICITY IN READING

David LaBerge
University of Minnesota

To begin my remarks, I would like to describe the general model which is a traditional way of representing information processing and the means of its control. Usually every process we wish to represent we put inside of a box. For example, when you present a stimulus, you put into one box the sensory feature detection, and a message from this box is then sent on to another box containing the pattern processor, and a message from this box is sent on to a semantic processing box, etc. Every time we want to talk about some process, we put it in a box. This way of modeling stems from the work of Shannon & Weaver (1949) and was developed further by Broadbent (1958).

Processing Within vs. Between Systems

In 1975, I tried to draw attention to the processing that goes on within a box and to suggest that it contrasts with the kind of processing that goes on between boxes. An example of within-box processing would be unitization of a perceptual pattern. Within this box we have feature detectors which are arranged in some hierarchical fashion leading to the arousal of activation in higher order perceptual codes, which are long-term memory structures. So, I agree with Buz that you come off the sensory surface into long-term memory right away. What it is that becomes automatic when you see a familiar pattern is some kind of an organization of these lower structures into the higher structures or codes. Lines representing organization of features into codes are not merely associations in the way that lines between patterns and their names are associations. That underestimates the richness of the properties that are going to be required to account for the organization which occurs within the perceptual system. At present, researchers are not in agreement as to what relationship the components make with the whole pattern, or even if there is such a thing as holistic processing. By the way, I should add that other processes such as semantic processes, are assumed by some investigators to feed down into these perceptual levels to help select what aspect of the total pattern is to determine the response in a given task. So this model is not always regarded as a pure bottom-up model.

Another kind of a between-systems processing that becomes automatic is the naming of patterns. Another example of within-system automaticity would be processing through of knowledge structures. Typically, lines within the semantic system are called links, which carry a good deal more than the properties of association which we've inherited from the British empiricists, namely facilitation and inhibition, and from neurological research, namely facilitation and inhibition. Another example of within-system processing which becomes automatic is motor control. Here, practicing a skill as if it were a sequence of associated responses may well be less effective than if it were practiced as a search for new ways to organize responses. If I may speculate for a moment, it seems clear to me that when I play an arpeggio on a piano, that I'm not going to get better by repeatedly articulating individual small units, but rather by exploring ways to position the hand as I group the notes.

We are now working with handshapes on keyboards to try to determine whether some of these organizational properties which we are borrowing from the work in perception can be generalized to the work in motor control.

Now, to contrast the organization-based automaticity within a system (or box) to the association-based automaticity between systems, we may consider the Shiffrin and Schneider (1977) theory, whereby letters or digits or other patterns become mapped onto some category. The category, I assume, can be represented as a semantic node, but I don't want to infer that this is a simple node, since it typically represents the complexities of meaning.

In summary, when automatic processing is within a box, it appears to obey different principles than when it occurs between boxes. However, I do not want to attempt to defend the assumption that these two categories of automaticity are the only ones. It is likely that people in this room can think of a third or fourth type of category.

One of the important consequences of separating two general modes of processing in this way is that the acquisition of automaticity may well be different for each case. I think that associating a visual code to its name is a matter of sheer practice, it's not a question of looking at the stimulus in different ways, or looking at the response differently. You simply exercise it. As Shiffrin said yesterday, exercising the mappings of visual things onto category responses produces automaticity. For example, training a subject to press the right button for a letter and the left button for a digit will show fast learning among college students, because of the kind of category learning that has already taken place before they enter the laboratory. On the other hand, the training of a person to produce a word unit, to consider those five letters now as one word, is not an exercise of simply looking at those letters again and again in some sequence. What probably happens is that the person begins to widen the integration extent for a unit as large as a word as opposed to the size of the integration for letters.

The learning of a larger motor unit, as I said before, is not a matter of repetition of finger sequences. Rather, I think you explore new hand shape feelings as you learn. As the Russian physiologist Bernstein (1967) said, practicing a skill is not a matter of repeating old solutions; it is a matter of problem solving. Practice session by practice session, what you're really looking for is ways that it can be done efficiently, ways of organizing such that you can give just one thought to many responses. In other words, you try to get the right kind of feel that's going to allow you to chunk as many little units as you can. In practicing a skill, you should not simply attempt to repeat old things, but you should attempt to look at what you are doing.

Unitization

I turn now to the problem of measuring unitization, a process which is assumed to take place within a system. We have tried to measure unitization, and to get an idea of what people are doing when they unitize. When you display a word, it is not a trivial matter to determine whether or not a person is looking at individual letters or the whole word. This problem is made more difficult because subjects have the capability of perceiving a word either way and, in fact, probably often do it either way, within the course of a daily college set of classes. One indicator of component vs. holistic word processing was described by Terry, Samuels and myself (1976). We presented words of different lengths, and we assumed that if a person were looking at a word in terms of components, whether it be a letter or spelling pattern, that the latency of classifying a word would increase with the number of components. We used words of lengths three to six and we asked subjects to make a go/no-go response on the basis of whether the words were names of animals or not. The main independent variable was word length, and the dependent variable was number of letter components. Another variable we considered at the same time was degradation of the letters. We conjectured that if subjects were looking at the component, then degrading it should affect its processing. We degraded individual letters by erasing sections of them, such that it took a person a little longer to identify them, but with enough time they could unambiguously identify each letter. We predicted that if a subject perceives familiar words in terms of the components, then there should be an increase in latency with word length. Degradation of the letters should then produce an interaction. So we predicted not only a significant slope of latency with word length for a normal presentation condition, but with the degradation condition a steeper slope. What we found was that for college students there is absolutely no effect of word length on the latency to identify the category of the word, even when letters were degraded.

We were not really satisfied with this result because we had no independent evidence that degradation was sufficiently strong to produce a clear test of the role of the letter quality. So, what we did was to induce the subjects to process component-wise, by presenting these same words in mirror image. The results showed significant slopes for the transformed non-degraded words and for the transformed degraded words, with the degraded condition showing the greater slope. So we tentatively concluded that the degradation was strong enough to affect letter processing when the subject was using letter components.

With this indicator of component vs. holistic processing in hand, we turned to the question of acquisition of unitary processing. We asked whether beginning readers would show component processing in this task when college students showed holistic processing. We took the task into the schools and tested subjects in second, fourth and sixth grades and compared them with college students. The results of this developmental study showed a large word length slope for second graders, with the slope decreasing as grade increased, until the slope reached zero for the college subjects. We tentatively concluded that word processing in early grades was by components (letters and/or spelling patterns), but that by college age, the processing was holistic.

Now, if we look at this thing called word length as a function of kind of task, we observe some interesting differences. So far we have described a categorization task in which we get a zero slope. However, in a physical matching task, one routinely gets word length effects with highly familiar words (Bruder, 1978; Eichelman, 1979). We have tried to eliminate the effect in the matching task by cuing words physically and semantically, and by using mismatching pairs that differ in all letters and positions but we still get a word length effect. From this kind of evidence, we conlcuded that the matching task requires a different kind of processing than a categorization task. Perhaps matching has more than just one stage. Perhaps there is unitary processing early followed by a check of the individual letters. What contrasts the matching task with the categorization task is the possibility that the subjects might perform the categorization task without processing components at any stage at all. Subjects might short-circuit what the component processing aspect of perceiving that is necessary in the matching task.

The lexical decision task (Butler & Hains, 1979) seems to fall somewhere between the matching and categorization tasks in this analysis. Again, I have heard of no consistent results which eliminate word length. Becker, in a recent manuscript, consistently gets word length effects across several manipulations. Naming a word seems to produce a slope effect with the number of components in many cases.

Now, taken together, the results which relate tasks to slope of the word length function seem to me to imply something about choice of training methods for word identification. Naming words as a training procedure does not transfer well to comprehension tests. Perhaps because you can sound out a word syllable by syllable, you stay at the component level. But for categorizing a word, you must unitize it before you can associate it to its proper meaning. You cannot get a meaning from letters, but only from the word. So, it would seem that asking a person to do a task that depends upon semantic processing of some sort can produce some kind of feedback which will encourage holistic processing. That would be one speculative conclusion relative to acquisition of unitization and eventual automaticity.

Automaticity

Now I turn to a measure of automaticity. We have been using a Stroop-type of test called the Flanker test, to get an estimate of how the processing of words may proceed automatically. Eriksen and Eriksen, in 1974, performed an experiment in the following way. They assigned four letters to two different responses, a lever press right or left. The letters S and C were assigned to one response and the letters H and K to the other response. They presented a display with these letters, such that one of these letters was in the center and three of one of the other four letters were in the flanking positions (e.g., CCCSCCC and HHHSHHH). Subjects were told to ignore the flanking letters. The warning signal was a fixation cross just below the middle letter. The main finding was that a display like CCCSCCC, which we call a compatible condition, is processed about 80 milliseconds faster than the display like HHHSHHH, which we call an incompatible condition.

I regarded this effect as an indicator of automatic processing of the flanker letters, and Shaffer and I attempted to repeat the procedure using words instead of letters. We used two buttons and we placed the words one on top of each in the following way:

PINE
IRON.
PINE

We used only one flanker item instead of three, because we found that the effect shows up with one flanker, and also Taylor found the effect using one flanker letter in a study which varied SOA in presenting the flankers. His results show a nice mapping of the flanker effect over time.

In the experiment by Shaffer and myself (1979), the word items allowed us to evaluate a category effect as well as a response effect. We assigned words from furniture and tree categories to one hand, and words from metals and clothing categories to the other hand. Thus, flanking words could come from not only a different category, but a different response relative to the target word. One other thing I might mention is that using categorization provides a real convenience in experiments of this sort because one can use a lot of different words; because categories are assigned to responses, one does not have to pretrain SR assignments or mappings of individual words. We put in neutral flankers occasionally in order to evaluate whether our compatibility/incompatibility effects were providing cost and/or benefits.

The results showed significant differences between neutral flanker conditions and compatible and incompatible flankers. In addition, it appeared that flankers of a different category but having the same response assignment, produced significantly slower responses than cases in which the flanker and target words came from the same category.

We can simplify the task just by using two categories, one assigned to each hand. We used words from the categories of body parts and animals. Our pilot work indicates that second graders show no compatible/incompatible effect, but sixth graders do show a compatible/incompatible difference, implying that, like the word length effect for unitization, second graders show no evidence of the flanker effect for automaticity. Unitization and automaticity seem to go together.

There is a problem in choosing the proper kind of neutral flanker to get an estimate of facilitation and interference in these flanker tasks. The proper choice of neutral flanker is not always an easy decision. You may devise a lot of different types, but not all of them fall in between the incompatible and compatible flankers. I'm not sure exactly why this variability occurs. But if any of you want to use the flanker task, and want to put in the neutral flankers in order to evaluate facilitation and inhibition, be careful in your choice of the kinds of neutral words.

Now, recall that we considered how one might unitize a pattern. I presume that what one does with word patterns is to look at more than just the letters, to order them, or consider them as a group. But also, one must consider a wider view and pick up perhaps whatever is happening inside the boundaries of the word. I don't like to use the term "relation", because it sounds very unscientific unless you're a mathematician, and we thought that we couldn't figure out what it was subjects were looking at in between the word boundaries mainly because one can't point to relations. Consider the difference between my two fingers; you can't point to it. You can point to my fingers, but not to the difference. As Bertrand Russell pointed out in his 1912 book on philosophical problems, consider that Edinburgh is north of London;

you can point to London and you can point to Edinburgh, but you never point to "north of". It simply doesn't exist; it's a Platonic universal, not something that is real. So it's no wonder that we have trouble communicating with each other about those things called relationships. But perhaps we could tell when the person is looking at something that must be other than just the little things you can point to or put nouns to.

The Attention Spotlight: An Attempt to Measure Its Size

What we did was the following. We presented words like "horse", a word with an "R" in the center, and "salad", a word with an "L" in the middle. If that middle letter was an "R", the subject was to press the right button; if the middle letter was an "L", the subject was to press the left button. Occasionally we presented an arrow pointing left or right, at the center or at the end letter position. Instructions were: "If you see an arrow appear and it points to the right, press the right button, if it points to the left, press the left button." This is not a very difficult task. We felt that if the response time to the arrows was sensitive to the size of the attentional spotlight focusing on the middle letter, then there would be a longer latency at the outside position of the arrows than at the center point. The other condition that we ran, in order to vary the spotlight size, was to ask the subject to classify the word as either an animal or a body part. We reasoned that in this case, if the person is looking at the word component-wise and reading from left to right, we should get a faster response at the left end position. On the other hand, if the word categorization condition induced subjects to focus on the word-as-whole, then the latency to the arrows should be constant across all three positions.

The results are based on 20% arrow probe trials with 12 subjects in each group. We found that the position of the arrow made a significant difference for the letter (R or L) condition but not for the word condition. In other words, the probe by task interaction was highly significant. Error rates were 1% in the letter task, 9% in the word task.

There is a problem with this procedure having to do with stimulus response compatibility. If the arrow is presented on the left and points to the left, the response is fast. But if an arrow in the left position points to the right, the response will be slower. We believed that it was very important to shift the response topography in such a way as to eliminate the compatibility effect if at all possible. So we repeated the experiment using a lever which we pressed forward in a go/no-go design, which seems to eliminate the problem. In this case, we used "Bs" against "Rs". For a "B" you press forward, and for an "R", you make no response. Occasionally we presented arrows, not only center and ends of the five-letter words, but on the outside too, because we wanted to map out more positions around the word pattern. There were five subjects in each group.

The same pattern of results appeared: a relatively flat curve for word judgments and a V-shaped curve for the letter judgments. Six of the subjects I was able to test across the two task conditions and thus get a better estimate of the relative height of the curves. The results showed that the center position gives about the same latency for the two conditions.

There was one other experiment which used this arrow test. This experiment separated the letters because it was thought that having a space might affect letter detection. When letters are closely spaced there may be a suppression effect much like a flanker test showed. So we put a space between each letter and we get about the same effect that we got before. In fact, you can almost put the curves over each other.

Here then is one way we're currently trying to understand what might be happening when a person looks at something as a unit, as opposed to looking at it as a component part.

References

Broadbent, D. E. *Perception and Communication*. London: Pergamon, 1958.

Bruder, G. A. Role of visual familiarity in the word superiority effects obtained with the simultaneous matching task. *Journal of Experimental Psychology: Human Perception and Performance*, 1978, *4*, 88-100.

Butler, B., & Hains, S. Individual differences in word recognition latency. *Memory and Cognition*, 1979, *87*, 68-76.

Eichelman, W. H. Familiarity effects in the simultaneous matching task. *Journal of Experimental Psychology*, 1970, *86*, 275-282.

Eriksen, B. A., & Eriksen, C. W. Effects of noise letters upon the identification of a target letter in a non-search task. *Perception and Psychophysics*, 1974, *16*, 143-149.

Shaffer, W. O., & LaBerge, D. Automatic semantic processing of unattended words. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 413-426.

Shannon, C. E., & Weaver, W. *The Mathematical Theory of Communication*, Univ. of Illinois Press, Urbana, 1949.

Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing, *Psychological Review*, 1977, *84*, 127-190.

Terry, P., Samuels, S. J., & LaBerge, D. The effects of letter degradation and letter spacing on word recognition. *Journal of Verbal Learning and Verbal Behavior*, 1976, *15*, 577-585.

A Test Between Two-State
and Continuous State Attention Models

Marilyn L. Shaw and Robert Mulligan

Rutgers - The State University of New Jersey

Lawrence D. Stone

Wagner & Associates

Abstract

Predictions of a class of two-state and a class of continuous state attention models are tested in an experiment requiring subjects to detect and locate targets. The model predictions are tested by comparing the effect of focused versus divided attention instructions on probability of correctly locating a target. The class of two-state models predict that target location performance will be better in the focused than in the divided attention condition while the continuous state models predict that the reverse will be true. Results from an experiment testing these predictions seem to favor the continuous state models.

In this paper we present an analysis of two-state versus continuous state attention models. Each class of models shares the assumption that at any given instant, the human observer has a limited amount of processing resources, capacity or attention. The classes differ, however, in assumptions concerning the internal representation of the stimulus used in deciding whether in a visual search task a target is present. In the class of two-state attention models, it is assumed that the internal representation can be treated as a two-valued random variable, e.g., this might be interpreted as a stimulus code in which a categorical code for the stimulus is retained instead of more detailed sensory information. This assumption, of course, does not exclude the possibility that the internal representation of the same stimulus is better treated as a more than two valued random variable at other points in time or other stages of processing. In the class of continuous state attention models it is assumed that the internal representation used in decision-making is best treated as a continuous valued random variable, e.g., this might be interpreted as a stimulus code in which a great deal of the sensory information is retained. Our analysis of these classes of models draws from three sources: work in psychology on attention to multiple sources of information (Kinchla, 1974; Eriksen & Spencer, 1969; Gardner, 1973; Shiffrin & Gardner, 1972; and Shaw, 1981); work in psychophysics on two-state and continuous state models (c.f. Krantz, 1969); and work in formal search theory (Stone, 1975; Mela, 1961; Kadane, 1971; Tognetti, 1968; Koopman, 1959)) concerned with performance optimization in detecting and locating targets. The outcome of our analysis is the derivation of testable predictions discriminating between two-state and continuous state attention models.

Organization of the Paper. We begin by presenting the class of search problems to which our attention models apply. This is followed by brief reviews of related literature in attention, psychophysics and search theory. Next, we give a description of the experimental paradigm we use and then present the formal models and their predictions.

THE SEARCH PROBLEM

Suppose there are n locations in space and the internal representation of information at each location can be treated as a random variable. Each random variable, x_{ki}, is one of two types: target (i=1) or nontarget (i=0) depending on whether a target or nontarget is present at location L_k. The X_{ki}'s may be dependent or independent. The pattern of types of random variables over n locations is denoted by $S_{i_1 i_2 i_3 \ldots i_n}$ where $i_k=0$ or 1 depending on whether a target or nontarget is in L_k. The probability of the pattern $S_{i_1 i_2 \ldots i_n}$ is denoted by $\Pi_{i_1 i_2 \ldots i_n}$. We consider two possible search objectives: detection and whereabouts. In a *detection search* the objective is to maximize the probability of correctly deciding whether a target is present or absent in at least one location regardless of where it is. In a *whereabouts search* the objective is to maximize the probability of correctly deciding which location, if any, contains the target. In these search problems, it is assumed that there is a limited resource (attention, processing capacity, search effort, search time) available for processing. One type of optimization problem is how to allocate the limited resource among the n locations to maximize one of the objectives given above. In this paper, we show how the optimal allocation of resources during search depends on the type of search and the information acquired during search. Furthermore, it is shown that the question of whether both optimization problems can be solved with the same allocation of resources depends upon assumptions about the information acquired during processing or search. More specifically, if the information acquired can be represented as a continuous valued variable (continuous state model) then simultaneous optimization of detection and whereabouts search is possible. If, however, the information acquired is represented in terms of a two-valued variable (two-state model), then each optimization problem requires a different allocation of resources. Our strategy is to use this difference in whether detection and whereabouts can be simultaneously optimized to provide an empirical test between these classes of models.

REVIEW OF RELATED WORK

Attention and Set Size Effects. It has been observed many times that dividing attention among more locations in space leads to a decrease in performance accuracy. Two possible explanations of this effect are that the decrease in accuracy may be due only to the increased opportunities for making an error, and that accuracy may also be lowered because less attention can be paid to each location. Set size effects in visual search have been found, for example, by Estes and Taylor (1964); Kinchla (1969, 1974(a), 1974(b) Eriksen and Spencer (1969); Gardner (1973); Shiffrin and Gardner (1972), to mention a few.

Recently, one of us (Shaw, Note 1) tested two specific attention models in several set size experiments. In these experiments, subjects were required to search through briefly flashed arrays of either two or four letters for a target letter. On each trial, the subject made two responses: one indicating "yes-no" whether a target was present or absent (detection search); and the other indicating which location, if any, contained the target (whereabouts search). In the case of two-letter arrays, subjects were instructed to attend equally to both letter positions and in four-letter arrays subjects were instructed to attend equally to the four letter positions. The two attention models tested were the Capacity Allocation Model and the Sample Size Model. The first is a two-state model and the second is a continuous state model. The two-state model is based upon the assumption that the conditional probability of detection of a target is the cumulative exponential function; the Sample Size Model is based upon the assumption that each internal random variable has a Gaussian probability function. In these models the amount of attention allocated to a location influences the parameters of these underlying probability distributions. The two non-attentional models assume that the internal random variables are continuous: one is based on the Gaussian probability function and the other on the exponential. These models do not assume that attention influences the parameters of the underlying probability functions and attention is not a parameter in the model. These models were compared to a boundary function for which it is assumed there are no processing resource limitations and no specific probability function is assumed. The predictions of the two attention models, two non-attentional models and the boundary function are presented in Figure 1 together with data from several experiments reported in Shaw (Note 1). In the figure, theoretical probability of a correct

290

Insert Figure 1 about here

---

location response for two locations is plotted against probability of a correct response for four locations for several models - the two attentional and non-attentional models and the boundary function. Data from individual subjects are indicated by the dots. The data clearly favor the attentional models, but the predictions of the Sample Size and Capacity Allocation Models are similar and so this data does not distinguish between them. Since both models give a good account of the data, yet differ considerably in their underlying assumptions, it is of interest to discriminate between them.

Psychophysics.

Two-State Threshold Models. Early investigators tested the sensitivity of sensory mechanisms in "threshold" experiments where the stimulus (target) was present on nearly all trials. A few "catch trials" (target not present) were inserted to insure that subjects did not say "yes" just because they knew that the target was always present. Relatively few catch trials were used because it was believed that a subject's tendency to report the perception of the target was generally influenced only by whether it was above threshold, and not by possible biases in response tendencies.

Classical high and low threshold models are based on the assumption that when a barely detectable stimulus is presented, the internal stimulus representation is in one of the two possible states, detected or not detected. In a high threshold model a detect state never occurs when the target is absent. If the target is present, the subject will be in a detect state if the energy of the stimulus is sufficient to evoke an internal state that exceeds some fixed value called a threshold. If the internal state exceeds the threshold, the observer reports the presence of the target; otherwise the response is "no." However, observers do sometimes respond "yes" when a target is absent. The presence of these false alarms led to the development of low threshold models in which a detect state may occur when the target is absent and an internal state exceeding the threshold does not necessarily lead to a "yes" response.

Luce (1963) presented a low threshold model which accounted for the false alarms clearly present in the data but not predicted by the high threshold model. He proposed that each exposure to a given stimulus had a fixed probability of producing a given internal state in the observer. If a target was present the observer would be either in a detect state with probability $\underline{p}$, or in a nondetect state with probability $1 - \underline{p}$. Similarly, when a nontarget was present the observer would be in a detect state with probability $\underline{q}$ or a nondetect state with probability $1 - \underline{q}$. Luce postulated two response-selection rules that generate pairs of hit and false alarm rates: (a) report target present whenever in the detect state and for a fraction $\underline{t}$ of the nondetect states (liberal strategy); or (b) report target absent whenever in a nondetect state and for a fraction $\underline{u}$ of the detect states (conservative strategy). By varying either $\underline{t}$ or $\underline{u}$ the observer could trade off hits and false alarms. Luce showed that the resulting ROC has two linear segments that meet where $\underline{t} = 0$ and $\underline{u} = 1$.

In this paper, the two-state model we present is a generalization of Luce's model from the single information source detection situation to situations where there is more than one source of information to attend to and there is a limitation on the total resource available for processing these sources. It should be noted, also, that our interpretation of the two-state model is broader than the traditional one. We assume the model applies whenever the stimulus code involves categorization of a stimulus into one of two possible states for the purpose of making a decision. Tests of the model do not address the question of whether this is the only possible code.

Continuous Models. In 1955, Swets, Tanner, and Birdsall questioned the assumption of a threshold that is independent of the subject's motivation. Following the Thurstonian tradition, they proposed that the set of possible internal states of an observer is best described in terms of a one-dimensional continuum of values. The internal representation of the stimulus is conceptualized as a random variable that can take on any value along this continuum. The probability that a stimulus evokes a value in a given interval along the continuum is specified by two overlapping probability distributions, one for the target stimulus (or signal) and one for the non-target stimulus (or noise). It is assumed that the subject selects some value along the continuum as the criterion, and classifies a stimulus as target or nontarget depending on whether the value evoked by it is above or below the criterion. This criterion, unlike the threshold, is under the subject's control and is influenced by motivational factors such as the relative frequency of targets and nontargets and the cost of making erroneous decisions. Errors arise because of the overlap of the target and nontarget probability distributions. A nontarget stimulus that evokes an internal representation above the criterion gives rise to an error traditionally called a false alarm; a target stimulus above the criterion leads to a correct response called a *hit*. Hit and false alarm rates vary depending on the observer's criterion.

In the present paper, we consider an extension of this model to situations with more than one source of information in which the target might appear and there is a limited resource for processing these sources. Two decision models have commonly been considered for the multiple information source detection paradigm (yes-no response): integration models and independent decisions models. In the former, it is assumed that the observation period results in values for the n random variables and the observer constructs a decision variable, y, as a weighted linear combination of these values:

$$y = w_1x_{1i} + w_2x_{2i} \ldots + w_nx_{ni}$$

If the decision variable y exceeds a criterion value, β, the observer says "yes;" otherwise she says "no." Independent decisions models, in contrast, assume a separate decision is made for each source before a final decision is made. Thus, for each source a decision variable, y_k, is constructed as follows:

$$y_k = 0 \quad \text{if} \quad x_{ki} < \beta_k$$

$$y_k = 1 \quad \text{if} \quad x_{ki} \geq \beta_k$$

The observer says "yes," target present, if $\Sigma y_k > c$. That is, when c or more of the decision variables, Y_k, indicate target present ($y_k = 1$).

There is evidence that subjects can use both kinds of decision strategies. However, for the class of search paradigms we are considering, the evidence overwhelmingly favors the independent decisions model (Shaw, Note 1; Mulligan and Shaw, 1981; Swennson and Judy 1981; Starr, Metz, Lusted, Goodenough, 1975). We test this model again and use it in our data analysis.

Optimization results for detection versus whereabouts search first appear in formal search theory (Koopman, 1956; Mela, 1961; Tognetti, 1968; Kadane, 1971; Stone, 1975). These investigations assume that search effort is discrete; that the target is always present in one of N locations; and that false detections of a target never occur. Resource allocation functions maximizing the probability of detecting a target under the above assumptions have been studied and reviewed by Kadane (1971) and Stone (1975). Here we illustrate the optimal allocation for detection search. The reader should consult Stone (1975) for a review of work on whereabouts search.

Suppose there is gold coin lost in one of N boxes. For each box, the prior probability for the coin is P(i). Let α_j be the probability of drawing the target in a single draw, given it is in Box j. Since draws are made with replacement, the probability of failing to draw the gold coin in the first k draws from Box j is

$$\beta(j,k) = \alpha_i \, (1 - \alpha_j)^{k-1}$$

Suppose also that a fixed number of draws K are permitted. We wish to allocate draws among the boxes in order to maximize the probability of drawing the gold coin given allocation ϕ; that is to maximize:

$$P[D] = \Sigma \, p(j)[1 - \alpha_j)^{\phi(j)}].$$

Note that

P(j) $\beta(j,k)$ is the discrete counterpart of the rate of marginal return. The optimal allocation places the next draw in the box with the highest rate of marginal return. In other words, the optimal allocation of draws is to make the n^{th} draw in cell i such that

$$P(i)\alpha_i\ (1-\alpha_i) = \max P(j)\,\alpha_i (1-\alpha_j)^{r(j,n-1,\phi)}$$

where $r(j,n-1,\phi)$ is the number of draws out of the first n that are placed in Box j by allocation plan ϕ.

Whereabouts search was first considered by Mela in the context of military search problems. Kadane (1969) and Stone (1975) have presented general results for this kind of search under the assumptions given above. Here we illustrate the optimal allocation for whereabouts search.

Now suppose that instead of trying to maximize the probability of drawing the gold coin in one of the K draws (detection search) we wish instead to maximize the probability of correctly specifying the box containing the target. Thus, if after K draws the gold coin is not found, we are allowed to guess which box contains the coin. In this case, how should the K draws be distributed among the boxes? The probability of correctly locating the target (correct whereabouts) is

$$P(CW) = p(i)[1-p(i)b(i,\phi(i))] + \Sigma p(j)b(j,\phi(j))$$

$$= p(i) + \Sigma\, p(i)b(j,\ \phi(i))$$

where Box i is the box that is guessed when the target is not detected or drawn. The optimal allocation of draws to maximize P(cw) is to choose Box i such that P(i) = max P(j). Furthermore, no draws are allocated to Box i and the K draws are allocated among the remaining N-1 Boxes so as to maximize the probability of drawing the coin if it is in one of these N-1 boxes. The reason for allocating no draws to box i is easily seen from the equation for P(CW): if the coin is in Box i, then it will never be found in one of the other N-1 boxes and Box i will be chosen. In this case, the gold coin will always be correctly located and so there is no point in wasting draws in Box i.

EXPERIMENTAL PARADIGM

The subject's task is to search an array of letters for a target letter. Arrays contain either two or four letters positioned at the corners of an imaginary square in the case of four letters and at the opposite ends of a diagonal in the case of two letters. There are four possible S_{ij} patterns for two letter displays and 16 possible S_{ijkl} patterns in the case of four letter displays.

There are two dependent variables. First, the probability the subject says "no," target absent for the detection search given pattern S_{ij} (S_{ijkl}) was presented. When displays contain two letters, the subject chooses either an upper or lower diagonal position for the whereabouts response. But, when the display contains four letters the subjects chooses either the positive or negative diagonal for the whereabouts response. On each trial the subject must make both a detection response and a whereabouts response, even when the detection response is "no."

There are two kinds of attention instructions for the four-letter displays: focused and divided attention.

Focused Attention. Here the subject is told to attend to the positive (negative) diagonal and ignore the other diagonal. For the detection response, the subject is instructed to say "yes" or "no" only to the information in the attended diagonal. For the whereabouts response, however, the subject is told to guess the unattended diagonal if the target is not detected in the attended diagonal.

Divided Attention. Here the subject is told to attend equally to all four locations. For the detection and the whereabouts response the subject is instructed to use information from all locations in making a choice.

THE FORMAL MODELS

Model assumptions can be divided into two groups: representation assumptions and decision-making assumptions. The former are concerned with the stimulus code (two-state versus continuous) used in decision-making and how attention affects this code. The decision-making assumptions are concerned with how a response is selected given the pattern of stimulus codes.

A Class of Two State Models

Representation Assumptions. First, attention is assumed to be a finite quantity, Φ. We denote by ϕ_k, the amount of attention allocated to L_k. It is assumed that Φ may be partitioned among locations and this does not change the total amount:

$$\Phi = \sum_{k=1}^{N} \phi_k .$$

This assumption was tested by Shaw and Shaw (1977).

The third assumption concerns the probability of a detection state for location L_k given the target is there ($X_{ki} = 1$) and ϕ_k has been allocated there: $b(\phi_k)$. It is assumed that $b(\phi_k)$ is continuous, concave and increasing

298

with ϕ_k. Naturally, $1 - b(\phi_k)$ is the probability of missing the target in location L_k. The Capacity Allocation Model makes the specific assumption that $b(\phi_k) = 1-e^{-\phi_k}$. The curve plotted in Figure 1 is based on this model.

The fourth assumption concerns the probability of a nondetection state for location L_k given a nontarget is there ($X_{k0} = 0$) and ϕ_k has been allocated there: $q(\phi_k)$. In the Capacity Allocation Model it was assumed $q(\phi_k) = 1$, that is, false alarms are assumed not to occur. More generally, we assume that $q(\phi_k)$ is continuous and increasing. The probability of a false alarm in L_k ($X_{k0} = 1$) is written as $1 - q(\phi_k)$.

Decision-Making Assumptions. Rather than describe a very general model, we specify a model for each of three experimental conditions: two letter displays, four letter displays with focused attention, and four letter displays with divided attention; and each response; detection and whereabouts. Furthermore, these response models assume that the target is equally likely to appear in each location or diagonal. A more general model is presented in Appendix 1.

Two Letter Displays - Detection Response. It is assumed that the subject may adopt either a liberal or a conservative strategy. If the subject adopts a liberal criterion, the response "yes" is given if a detect state occurs in either or both locations. In addition a "yes" response is assumed to occur with probability $\underline{u}$ when neither location results in a detect state. If the subject adopts a conservative criterion, then a detect state in one or more locations is assumed to produce the response "yes" with probability $\underline{t}$, and "no" with probability $1-\underline{t}$. The response is always "no" if neither location results in a detect state.

299

Two-Letter Displays - Whereabouts Response. If only one location results in a detect response then that location is selected for the whereabouts response. If both locations or neither location results in a detect state, then one of the locations is chosen at random.

Four-Letter Displays - Focused Attention. Since S is attending to only two locations the detection response model is the same as for the two-letter displays. If either location in the attended diagonal results in a detect state then the attended diagonal is chosen for the whereabouts response; otherwise, the attended diagonal is selected.

Four-Letter Displays - Divided Attention. For the detection response the subject may adopt a liberal or conservative strategy. In an extreme liberal strategy the subject may say "yes" to some fraction of the trials where there are no detect states, while in an extreme conservative strategy, the subject may say "no" to some fraction of trials where there are four detect states.

For the whereabouts response the subject is assumed to choose the diagonal with the greatest number of detect states or to pick randomly when the diagonals have the same number of detect states.

Two-State Model Predictions. Let $P(CW/FA, S_{ij})$ and $P(CW/DA, S_{ij})$ be the probability of a correct whereabouts response in the focused attention and divided attention conditions, respectively. Similarly, we denote these probabilities for the detection response by $P(CD/FA, S_{ij})$ and $P(CD/DA, S_{ij})$.

Predictions for Whereabouts Response. The two-state attention model predicts that the probability of a correct whereabouts response will be greater in the focused attention condition than the divided attention

300

condition. This is analogous to the theorem we described from search theory. However, the present model assumptions differ somewhat from the assumptions made in formal search theory. First, we assume that it is possible for a detect state to occur when a nontarget is in a location. Second, we assume search effort or capacity is a continuous quantity. Previous investigations of the whereabouts problems in search theory have assumed that a detect state can only occur when the target is present and that search effort is discrete. In Appendix 1 we show that for the case of two locations and a uniform prior that on the average, the probability of a correct whereabouts response will be greater for the focused attention condition than for the divided attention condition. That is,

$$P(CW/FA) > P(CW/DA). \tag{1}$$

If one does not assume a uniform prior, then whether focused or divided attention is the superior strategy depends in a somewhat complicated way upon the relation between $b(\phi)$ and $q(\phi)$.

The inequality in (1) is meaningful when only one target is present in the four-letter displays or when two targets are present in the same diagonal. If a target appears in both diagonals, then the subject is correct no matter what the choice; conversely, if no targets are present then no choice of location can be correct.

Whether the prediction in the above inequality can be observed depends on the size of the difference between these two probabilities and the precision of the data collected. We computed the expected difference in P(W/FA) and P(CW/DA) for the Capacity Allocation Model as a function of the difference between the false alarm rate when attention is focused versus when attention is divided. This difference $[P(CW/FA) - P(CW/DA)]$ is plotted in Figure 2 for two different values of Φ. The predicted difference in the figure ranges from .05 to .45. Thus, it appears that one can ensure that the expected difference for this model can be made sufficiently large by using display energy conditions where the estimated Φ corresponds to predictions of $P(CW/FA) - P(CW/DA) > .10$.

Predictions for the Detection Response. The value of the inequality (1) depends completely on whether subjects can and do follow the instructions to attend to a single diagonal in the focused attention conditions. To check this we compare detection performance in the focused attention condition to detection performance in the two-letter array condition. Since in both cases the subject need only process two locations, we should find that detection performance is the same in both cases:

$$P(CD/FA, \text{two diagonals}) = P(CD/DA, \text{one diagonal}). \tag{2}$$

A Class of Continuous State Attention Models

Representation Assumptions. First, attention is assumed to be a finite quantity, Φ. Second, if ϕ_k is the amount allocated to L_k, it is assumed that

$$\Phi = \sum_{k=1}^{N} \phi_k .$$

That is, dividing attention does not result in any change in the total amount. Let $f_t(X)$ and $f_d(X)$ denote the probability density function of the random variables X_{k1} and X_{k0} associated with L_k when target present and absent, respectively. It is assumed that f_t and f_d are symmetric, unimodal distributions and have the same functional form. For convenience we assume that the mean of the target random variable is greater than the mean of the nontarget random variable. Attention is assumed to influence the variability (noise level) of the internal representation. Specifically, we assume $\sigma^2_{x_{ki}} = \frac{\sigma^2}{\phi_k}$ where σ^2 is constant representing the variance inherent in the stimulus representation not influenced by attention. Evidence that this last assumption is consistent with data from set size studies is presented in Figure 1. An attention model very similar to this model has been studied by Luce and Green (c.f. Green and Luce, 1973; Luce, 1977) in auditory detection and magnitude estimation paradigms.

Decision-Making Assumptions.

Two-Letter Displays - Detection Response.

In the same experimental paradigm we use here, Shaw (Note 1) found that the detection responses of subjects were consistent with the independent decisions model and not with the integration model. This model assumes that the probability of a "no" response given S_{ij} is

$$P(no/S_{ij}) = P(X_{1i} < \beta_1)P(X_{2j} < \beta_2).$$

Two-Letter Displays - Whereabouts Response.

Most models of the continuous variety assume that when forced to choose between n random variables the subject chooses the response whose associated random variable is the maximum (minimum) of the n. We likewise make this assumption and so the probability location L_k is chosen given S_{ij} is given by

$$P(no/S_{ij}) = P(X_k = \max(X_{1i}, X_{2j}))$$

Four Letter Displays - Focused Attention. Since the subject is attending to only two locations, the decision model for the detection response is the same as for the two letter displays. For the whereabouts response the subject is assumed to choose the attended diagonal whenever one of the associated random variables exceeds its criterion; otherwise the subject chooses the unattended diagonal.

Four Letter Displays - Divided Attention. For the detection response the subject is assumed to say "yes" whenever at least one random variable exceeds its criterion; otherwise the subject says "no." For the whereabouts response we let $r = \max (X_{1i}, X_{2j})$ denote the maximum of the random variable associated with the positive diagonal and $s = \max (X_{3k}, X_{4l})$ denote the same quantity for the negative diagonal. It is assumed that the subject chooses the positive diagonal if $r = \max (r,s)$ and the negative diagonal if $s = \max (r,s)$.

Continuous State Model Predictions.

Whereabouts Response. In contrast with the two-state model, the continuous state models make the prediction that the same allocation of attention optimizes both whereabouts probability correct and detection probability correct (see Appendix 1). This predicts that subjects should perform no better or more poorly in whereabouts when they are in the focused attention conditions (ignoring a location to be guessed). Thus, we have

$$P(CW/DA, S_{ij}) \geq P(CW/FA, S_{ij}). \qquad (3)$$

Detection Response. Again, the value of the inequality (3) depends completely on whether subjects can and do follow the instructions to attend to a single diagonal in the focused attention conditions. To check this assumption we compare detection performance in the focused attention condition to that in the two-letter array condition. Since in both cases the subject need only process two locations, we should find that detection performance is the same in both cases:

$$P(\text{CD/FA two diagonals}) = P(\text{CD/DA one diagonal}) \qquad (4)$$

The Continuous State Model and its predictions are presented in more detail in Appendix 1.

SUMMARY

Two classes of attention models make contrasting predictions for the probability of a correct whereabouts response for the different attention instructions we use. The two-state models predict that the whereabouts response will be superior in the focused attention condition whereas the continuous state models predict that performance will be no better or worse in the focused attention condition when compared to the divided attention condition. Data from the detection response is simply used to confirm the assumption that subjects do indeed ignore the unattended diagonal when so instructed.

For our experiment we choose the two-location case where targets are equally likely in all locations and locations are independent. Though the theorems we draw upon apply to the dependent case, this does not cause a problem for us since we tell our subjects the optimal allocation strategy for the dependent case $(p(1) = 1-p(2))$ and then analyze probability of a correct whereabouts only on those trials on which the target is in only one location or the other.

METHOD

Subjects. Four adults, three females and one male, having normal or corrected to normal vision, served as observers. They were paid $3.00 per hour.

Stimuli. The stimulus displays were similar to those used previously by Shaw (Experiment 4). Both two-location and four-location displays were used. The stimulus at each location consists of a single letter flanked by dollar signs. In the two-location display, the three-character strings appear at the ends of an imaginary diagonal, separated by 1.25 degrees of visual angle. In half of these displays the characters appeared at the opposite ends of the positive diagonal (at the lower-left and upper-right of the screen); in the other half, they appeared at opposite ends of the negative diagonal (upper-left, lower-right). Four-location displays contained character strings at all four positions described by the ends of the two intersecting diagonals.

Varied mapping (Schneider & Schiffrin, 1978) was employed, i.e., the target letter changed from trial to trial. Both targets and nontargets were randomly selected on each trial from the following group: F, H, J, K, L, M, N, T, V, W, X, Y, and Z. Each letter was used equally often as target and nontarget and each appeared equally often in each position.

The a priori stimulus probability distribution for the two single diagonal conditions was as follows: $P(S_{00}) = .40$, $P(S_{01}) = .25$, $P(S_{10}) = .25$, and $P(S_{11}) = .10$.

Stimuli in the four-location displays are denoted by S_{ijkl}, where i and j denote the event in the upper and lower positions of the positive diagonal, and $\underline{k}$ and $\underline{l}$ the events in the upper and lower positions of the negative diagonal, respectively. The stimulus probabilities approximate independence and in this condition were: $\underline{P}(\underline{S}_{0000}) = .40$, $\underline{P}(\underline{S}_{0001}) = .13$, $\underline{P}(\underline{S}_{0010}) = .13$, $\underline{P}(\underline{S}_{0100}) = .13$, $\underline{P}(\underline{S}_{1000}) = .13$, $\underline{P}(\underline{S}_{0101}) = .02$, $\underline{P}(\underline{S}_{0110}) = .02$, $\underline{P}(S_{1001}) = .02$, and $\underline{P}(\underline{S}_{1010}) = .02$. Thus, the probability of a target on only the negative diagonal was .26, on only the positive diagonal, .26, and on both diagonals, .12. Note that two targets never appeared on the same diagonal and stimulus patterns with 3 and 4 targets did not appear. They occur with such small frequency in the independent case that we did not include them.

Procedure. Subjects viewed the displays binocularly through a Tektronix viewing hood mounted 87 cm from the front surface of a Tektronics Model 610 cathode ray tube. The CRT was under the control of a DEC PDP-12 computer, using its standard letter format.

Each trial began with a 2-sec presentation of the letter chosen as target, followed by a 3-sec fixation period, and then the 10-msec stimulus display. The fixation patterns consisted of two or four dots corresponding to the display locations, and an additional dot at the center of the display, i.e., at the point of intersection of the two diagonals. The subjects were instructed to prepare for each display presentation by fixating on the center dot. The importance of maintaining consistent fixation on all trials in all session types was stressed. Display presentation was followed by a 2-sec period in which a "Yes-No" detection response was made, a 2-sec period in which a location judgment was made, and a 2-sec feedback period during which the target letter and stimulus pattern were redisplayed.

On two-location trials, the location judgment required a forced choice between the two stimulus locations. On four-location trials, the observer's task was to specify in which diagonal -- positive or negative -- the target had appeared.

In both conditions, subjects were required to make a location judgment regardless of the outcome of their detection judgment.

Stimulus displays were presented in three types of sessions, corresponding to the three attentional strategies subjects were instructed to follow. In a divided attention session (DA) subjects were instructed to divide their attention equally and simultaneously between all display locations, so as to maximize correct detection responses. In a focused attention positive diagonal (FAP) session, subjects were required to attend only to the positive diagonal and to ignore the other two locations. They were told to base their Yes-No response solely on the attended diagonal. For the location question, they were instructed to choose the positive diagonal if their detection response was positive, and otherwise, guess the negative diagonal. Similar instructions were given for the focused attention on negative diagonal sessions (FAN) except subjects were to base their responses on information from the negative diagonal.

Each DA session consisted of 20 unrecorded practice trials followed by three blocks of 152 trials -- a positive diagonal two-location block, a negative diagonal two-location block, and a four-location block. Order of blocks within sessions was counterbalanced. Subjects received ten-minute breaks between blocks of trials.

Each FA session was comprised of 50 unrecorded two-location trials, followed by two blocks of 152 four-location trials. Subjects followed the same instructions (FAP or FAN) for both blocks within a session. The two-location trials, always presented first, were given to help the subjects focus their attention on the appropriate diagonal for that session. Subjects reported this to be a useful procedure in helping them to follow whereabouts instructions. As an additional aid, the fixation pattern on four-location trials indicated only the diagonal to which the subject should attend.

Data was collected in 17 sessions (9DA, 4FAP, and 4 FAN), with whereabouts sessions alternated with detection sessions. In the 9 DA sessions a total of 2736 two-location trials (1368 for each diagonal) and 1368 four-location trials were presented. In the 8 focused attention sessions, 1216 FAP and 1216 FAN trials were presented.

Prior to data collection all subjects had participated in eleven practice sessions -- 6 DA, 2 FAP, and 2 FAN -- a minimum of 608 trials per condition. During the first two sessions, both of which were DA sessions, display durations were gradually decreased from 100 to 10 msec. Thereafter, detection and whereabouts sessions were alternated. During these practice sessions, display brightness was differentially adjusted for each subject to achieve the desired accuracy levels.

RESULTS

Response probabilities (sample sizes in parentheses) are presented in Table 1 for all subjects and experimental conditions. Standard deviations for these estimated probabilities ranged from .03 to .009 depending on the value of the estimate and the number of observations.

Insert Table 1 About Here

Set Size Effects. In Figure 3 the "set size effect" in this experiment is graphed as the probability of a correct whereabouts response for the one diagonal condition (averaged over positive and negative diagonal data) versus this probability for the two diagonal, divided attention condition. For each subject, the data are consistent with the two attention models and are well below the boundary function. The theoretical functions in Figure 2 differ from those found in Figure 1 because in our experiment subjects must choose between diagonals in the four location condition. For this reason, the boundary curve was computed using a result from Green and Weber (1980) and used in Shaw (Note 1) in analyzing data from a similar experiment. The equation used is

Insert Figure 3 About Here

$$P_4 = 1 - \frac{2}{3}\left[1 - \frac{P_2}{P_2 + 3(1-P_2)}\right],$$

where P_4 is the probability of a correct whereabouts response given four locations and P_2 is the probability of a correct whereabouts response given two locations. Each estimated probability on the x - axis is based on about 1000 observations and those on the y- axis each is based on about 720 observations. These estimates were obtained by pooling data over occurrence of the target in the different possible positions.

Independent Decisions. The "yes-no" data from the two location (one diagonal) conditions were tested against the prediction that

$$P(\text{No} \mid S_{ij}) = P(X_{1i} < \beta_1)\, P(X_{2j} < \beta_2);$$

using a statistic derived in Shaw (1980). This test uses the fact that

$$\ln P(\text{No} \mid S_{ij}) = \ln P(X_{1i} < \beta_1) + \ln P(X_{2j} < \beta_2).$$

The standard normal z-scores of this test are presented in Table 2. The data of three of the four subjects is consistent with the independent decisions model. These results are also consistent with the assumptions that 1) subjects divided attention equally between the locations; and 2) the chance of a detect state given a target ($b(\phi)$) is the same as the chance of a nondetect state given a nontarget ($q(\phi)$). Evidence favoring this assumption is important for the analysis of the whereabouts response data testing the predictions of the two classes of models.

Probability Of A Correct Whereabouts Response. In Figure 4 the probability of a correct whereabouts response for the two diagonal divided attention condition is plotted versus this probability for the focused attention condition.

Insert Figure 4 About Here

Recall that the class of continuous state models predicts that P(CW | DA) $\geq$ P(CW | FA); this means that points in the graph should lie on or below and to the right of the diagonal. The diagonal line represents the predictions for the continuous state model and the choice of the optimal criteria. On the other hand, the class of two-state models predict that P(CW | DA)< P(CW | FA); this means data points should lie above and to the left of the diagonal line. The predicted relationship between P(CW/DA) and P(CW/FA) for the Capacity Allocation Model and the optimal decision rule are represented by the dashed line. These predictions were based on the assumption that the conditional probability of a correct detect state is the same as for a correct nondetect state and this probability is given by the cumulative exponential function: $1 - e^{-\phi}$. If the probability of a nondetect state given a distractor is smaller than the probability of a detect state given a target then this dashed line will move closer to the diagonal. In fact, if the chance of a nondetect state given a distractor is very small then the dashed line will fall well below the diagonal. This point is discussed in Appendix 1 under two state models. The "yes-no" data presented in Table 1 are consistent with the assumption that the probability of a nondetect state given a distractor ($q(\phi)$) is greater than or equal to the probability of a detect state given a target ($b(\phi)$). (See Table 2).

In summary, it seems that the prediction of the class of continuous state models is supported by this analysis of the data.

Focused Versus Divided Attention. The validity of conclusions reached from the data in Figure 4 depends upon whether subjects indeed focus attention on a single diagonal and ignore the information in the unattended diagonal. We now consider for each class of models the consequences of a failure of subjects to focus attention when instructed to do so. Failure to focus attention may result in dividing attention either equally or unequally between diagonals and then using the same decision strategy as in the divided attention condition. For the class of continuous state models this will result in points that will lie on or below the diagonal (on the diagonal if attention is divided equally and the optimal criteria is used). The same implication holds for the class of two state models. Thus, failure of subjects to follow this instruction will probably result in data below the diagonal and consistent with the predictions of the class of continuous state models. We can test the following consequences of focusing attention and ignoring the unattended diagonal for the class of two state models. These predictions do not necessarily hold for the class of continuous state models because subjects may shift criteria from condition to condition. However, since it is only the class of two state models that are of concern in this prediction,this does not affect the generality of our conclusions. First, it is predicted that the probability of a no response given only distractors in a diagonal should be the same whether that diagonal is the only one present or simply the attended diagonal. This should be true independent of whether a target is present or absent in the unattended diagonal (Tests 1 and 2 in Table 3). Second, the probability of a "no" response given one target present should be the same whether that diagonal is the only one present or simply the attended diagonal (Test 3).

Insert Table 3 About Here

In Table 3 we present the data relevant to these predictions and the results of testing them. For these tests the estimated probabilities were arrived at by pooling over positive and negative diagonal data. Inspection of the table reveals that only the data of subjects 2 and 4 favored the conclusion that they ignored the information in the unattended diagonal when they were instructed to do so.

DISCUSSION

There are four important results to be summarized from our data. First, we found the effect of set size under divided attention conditions is consistent with attention models as opposed to nonattentional models (see Figure 3); the data is also consistent with previous set size findings (see Figure 1). In addition, a good account of the set size effects in ours and other data is given by two specific attention models: the Capacity Allocation Model and the Sample Size Model. The former is a member of the class of two-state models we have presented while the latter is a member of the class of continuous state models. Second, an analysis of the one diagonal detection data indicated that 1) subjects appear to use an independent decisions rule; and 2) that if the two-state model holds, then our data are consistent with the assumption that the probability of a detect state given a target is the same as the probability of a nondetect state given a nontarget $[(b(\phi) = q(\phi)]$. Third, an analysis of detection data for four-location divided attention versus focused attention revealed that only two of our four subjects seemed able to effectively ignore information in the diagonal to be ignored. Finally, the data of these two subjects were consistent with the prediction of the class of continuous state models rather than the class of two-state models.

Our predictions discriminating between two-state and continuous state attention models may be applied to information processing situations other than visual search. For example, consider detection and location of an auditory signal; detection and identification of pitches; and detection and identification of spatial frequencies. However, for these results to apply the situation must involve a limitation on processing resources and control over the allocation of this resource. In our experiments, resource allocation was accomplished by instructing to focus or divide attention. Alternatively, the experimenter may overtly control the subject's allocation of attention. For example, Shaw and Bates (Note 2) tested subjects in a visual search experiment in which only two fixations were allowed in searching for a target. When subjects were asked to optimize a whereabouts response, they allocated all their fixations to one area and guessed the ignored area when a target was not detected. In the present experiments a similar overt control of attentional allocation could have been achieved by not displaying the ignored diagonal, but still allowing it as a whereabouts choice. Other constraints on the application of our model predictions are: 1) stimuli must be classifiable into one of two types (e.g. target and nontarget); 2) it must be possible to experimentally approximate equal detectability of each stimulus type for all sources of information; and 3) it must be possible to put multiple sources of information into two groups and increase the division of attention by increasing the number of subcomponents in each group. In our experiment this last constraint was satisfied by defining the areas to be attended as the diagonals of a square. This had another advantage over the use of only two positions in space. In focusing attention on a diagonal there is no advantage to fixating at the upper versus lower position as would be the case if subjects were instructed to focus attention on an upper versus lower position.

The classes of two-state and continuous state attention models we have considered share two basic types of assumptions. First is the limitation on capacity and its conservation over different allocations. Second is the assumption that subjects make separate decision for each location and then pool these decisions for the "yes-no" response. Thus, in each model the information about each location is coded categorically. The models differ in whether the criterion applied to internal representations can be varied: for the two-state model the category boundary or criterion is fixed while for the continuous state model it can vary. Thus, the two-state models are applicable whenever the stimulus domain seems to have relatively fixed category boundaries while continuous state models are applicable whenever category boundaries in a stimulus domain are variable. From this view the results of the present experiment suggest that letters seen under the degraded viewing conditions typical of tachistoscopic studies invoke representations the interpretations of which can be as variable and subject to motivational factors as, for example, an auditory signal embedded in noise.

REFERENCE NOTES

Shaw, M. Attending to Multiple Sources of Information. Manuscript submitted to Cognitive Psychology.

Shaw, M.L. & Bates, M. On the Difference Between Detecting and Locating Targets. Paper presented at the meeting of the Mathematical Psychology Society, Providence, Rhode Island, August 1979.

REFERENCES

Eriksen, C.W., & Spencer, T. Rate of information processing in visual perception: Some results and methodological considerations. Journal of Experimental Psychology Monograph, 1969, 79 (2, Part 2).

Estes, W.K., & Taylor, H.A. A detection method and probabilistic models for assessing information processing from brief visual displays. Proceedings of the National Academy of Sciences, 1964, 52, 446-454.

Gardner, G.T. Evidence for independent parallel channels in tachistoscopic perception. Cognitive Psychology, 1973, 4, 130-155.

Green, D.M., & Luce, R.D., Speed-accuracy tradeoff in auditory detection. In S. Kornblum (Ed.), Attention and Performance IV, New York: Academic Press, 1973.

Green, D.M. & Weber, D.L. Detection of temporally uncertain signals. Journal of the Acoustical Society of America, 1980, 67 (4), 1304-1311.

Judy, P.F., Swensson, R.G., Massimiliano, S. Lesion detection and signal-to-noise ratio in CT images. Medical Physics, 1981, (in press).

Kadane, J.B. Optimal Whereabouts Search. Operations Research, 1971, 19, 894-904.

Kinchla, R.A. Temporal and channel uncertainty in detection: A multiple observation analysis. Perception & Psychophysics, 1969, 5, 129-136.

Kinchla, R.A., & Collyer, C.E. Detecting a target letter in briefly presented arrays: A confidence rating analysis in terms of a weighted additive effects model. Perception and Psychophysics, 1974, 16, 117-122.

Kinchla, R.A. Detecting target elements in multi-element arrays: A confusability model. Perception & Psychophysics, 1974, 15, 149-159.

Koopman, B.O. The theory of search: Pt. III: The optimum distribution of search effort. Operations Research. 1957, 5, 613-626.

Krantz, D.H. Threshold theories of signal detection. Psychological Review, 1969, 76, 308-324.

Luce, R.D. Thurstone's discriminal processes fifty years later. Psychometrika, 1977, 42, 461-489.

Luce, R.D. A threshold model for simple detection experiments. Psychological Review, 1963, 70, 61-79.

Mela, D.F. Information theory and search theory as special cases of decision theory. Operations Research, 1961, 9, 907-909.

Mulligan, R. & Shaw, M. Multimodal signal detection: Independent decisions vs. integration. Perception & Psychophysics, 1981, 28(5), 471-478.

Shaw, M.L. Identifying attentional and decision-making components in information processing. In R.S. Nickerson (Ed.), Attention and Performance VIII. Hillsdale, NJ: Erlbaum, 1981.

Shaw, M. Reaction time in reading a tachistoscopic display for a memory set item. Perception & Psychophysics, 1977, 21(1), 15-25.

Shaw, M.L. & Shaw, P. Optimal allocation of cognitive resources to spatial locations. Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 201-211.

Shiffrin, R.M., & Gardner, G.T. Visual processing capacity and attentional control. Journal of Experimental Psychology, 1972, 93, 72-83.

Starr, S.J., Metz, C.E., Lusted, L.B., Goodenough, D.J. Visual detection and localization of radiographic images. Radiology, 1975, 116, 533-538.

Stone, L.D. Theory of Optimal Search. New York: Academic Press, 1975.

Tognetti, K.P. An optimal strategy for a whereabouts search. Operations Research, 1968, 16, 209-211.

Footnotes

* This research was sponsored by a National Science Foundation Grant, BNS 77-26296, to Rutgers--The State University. Requests for reprints should be addressed to the first author.

TABLE 1

RESPONSE PROBABILITIES

Two Locations

Positive Diagonal

SUBJECT	Probability of a "No" Response				Probability of a correct whereabouts	
	P_{00}	P_{10}	P_{01}	P_{11}	Upper	Lower
1	.90(392)	.09(256)	.12(252)	.01(106)	.96(256)	.95(252)
2	.85(419)	.16(264)	.15(263)	.03(112)	.95(264)	.96(263)
3	.89(417)	.11(266)	.19(266)	.08(112)	.96(266)	.90(266)
4	.85(410)	.17(260)	.17(260)	.04(104)	.95(260)	.94(260)

Negative Diagonal

SUBJECT	Probability of a "No" Response				Probability of a correct whereabouts	
	P_{00}	P_{10}	P_{01}	P_{11}	Upper	Lower
1	.88(372)	.18(251)	.11(250)	.03(107)	.91(251)	.96(250)
2	.89(418)	.19(265)	.11(265)	.05(112)	.97(265)	.93(265)
3	.87(418)	.08(266)	.11(266)	.04(112)	.93(266)	.97(266)
4	.86(418)	.19(266)	.18(266)	.05(112)	. 94(266)	.95(266)

TABLE 1 (Cont'd)

Four Locations

Divided Attention

	Probability of a "No" Response			Probability of a Correct Whereabouts Response	
		One Target Present*			
SUBJECT	No Target	Positive Diagonal	Negative Diagonal	Positive Diagonal	Negative Diagonal
		(one target present)			
1	.810(509)	.160(341)	.240(340)	.870(341)	.800(340)
2	.750(528)	.225(357)	.140(357)	.870(357)	.870(357)
3	.770(533)	.250(359)	.270(360)	.880(359)	.845(360)
4	.760(534)	.241(360)	.260(360)	.860(360)	.851(360)

Focused Attention

	Probability of a "No" Response			Probability of a Correct Whereabouts Response	
		One Target Present			
SUBJECT	No Target	Positive Diagonal	Negative Diagonal	Positive Diagonal	Negative Diagonal
1	.890(480)	.185(319)	.570(319)	.850(319)	.840(319)
2	.870(475)	.150(318)	.875(318)	.780(318)	.885(318)
3	.830(480)	.300(319)	.560(316)	.750(319)	.890(316)
4	.850(470)	.165(310)	.880(309)	.801(310)	.825(309)

Focused Attention

	Probability of a "No" Response			Probability of a Correct Whereabouts Response	
		One Target Present			
SUBJECT	No Target	Positive Diagonal	Negative Diagonal	Positive Diagonal	Negative Diagonal
1	.870(480)	.500(320)	.235(318)	.835(320)	.830(318)
2	.850(452)	.870(318)	.170(318)	.905(318)	.820(318)
3	.880(480)	.705(319)	.235(320)	.890(319)	.755(320)
4	.865(470)	.875(309)	.165(310)	.860(309)	.850(310)

*There is very little data for cases in which two targets were present so we have not included these figures.

TABLE 2

Standard Normal Scores* for
Tests of the Independent Decisons Model

Two Locations

SUBJECT	Positive Diagonal	Negative Diagonal
1	.18	-.48
2	-.11	1.625
3	-3.17	-2.56
4	-.31	-.51

*The prediction of the independent decisions model is that

$$\ln P_{ij} = \ln P(X_{1i} < \beta_1) + \ln P(X_{2j} < \beta_2).$$

This means $P_{11} + P_{00} = P_{10} + P_{01}$. The test statistic we use is

$$Z = \frac{\ln P_{00} + \ln P_{11} - \ln P_{10} - \ln P_{01}}{S}$$

where

$$S = \sqrt{\frac{(1-P_{ij})}{N_{ij}P_{ij}}}.$$

TABLE 3

Three Tests of Whether Subjects Ignored the Unattended Diagonal

TEST 1

	ONE DIAGONAL	TWO DIAGONALS (FOCUSED ATTENTION)	
	$P(NO \mid S_{00})$	$P(NO \mid S_{10}, S_{01}$ in unattended diagonal and S_{00} in attended diagonal)	Z
SUBJECT			
1	.890(764)	.535(639)	14.89
2	.870(837)	.872(636)	- .11
3	.880(835)	.630(635)	11.32
4	.855(828)	.877(618)	- .92

TEST 2

	ONE DIAGONAL	TWO DIAGONALS (FOCUSED ATTENTION)	
	$P(NO \mid S_{00})$	$P(NO \mid S_{00}$ in attended and unattended diagonal)	Z
SUBJECT			
1	.890(764)	.880(960)	.64
2	.870(837)	.870(927)	.02
3	.880(835)	.850(960)	1.86
4	.855(828)	.860(940)	- .27

TEST 3

	ONE DIAGONAL	TWO DIAGONALS (FOCUSED ATTENTION)	
	$P(NO \mid S_{10}, S_{01})$	$P(NO \mid S_{10}, S_{01}$ in attended diagonal and S_{00} in unattended diagonal)	Z
SUBJECT			
1	.125(1009)	.270(637)	-7.49
2	.175(1057)	.160(636)	.78
3	.123(1064)	.268(639)	-7.56
4	.178(1052)	.165(620)	4.69

TABLE 4

CONTINUOUS STATE MODEL

Examples from Monte Carlo Study

PARAMETERS					PROBABILITY OF CORRECT WHEREABOUTS	
μ	σ	β	η	Φ	P(CW/DS)	P(CW/WS)
2	1	1	2	4	.977	.977
			3	4	.910	.891
			4	8	.942	.919
			6	12	.915	.888
			8	8	.711	.693
			8	16	.893	.866
			8	32	.986	.969
1	1	5	4	4	.552	.540
			4	8	.678	.660
			4	16	.883	.799
			4	32	.942	.919
			4	64	.993	.984
1	1	1	2	4	.841	.739
			4	4	.552	.528
1	1	.5	2	4	.840	.840
			4	4	.550	.540
			6	10	.540	.540
1	1	.75	2	4	.841	.812
			4	4	.552	.536
			6	10	.546	.527

Appendix 1

TWO STATE MODEL

Notation

Symbol	
L_k	Location k.
X_{ki}	Internal Representation (code) of the stimulus in L_k where $i = 0$ for nontarget and $i = 1$ for a target.
S_{ij}	Stimulus Pattern: $i,j = 0,1$ depending on whether nontarget or target in L_k.
$P(L_k)$	Probability the target appears in L_k.
ϕ	The total amount of attention, capacity, processing resources or search effort.
ϕ_k	The amount of attention allocated to L_k.
$b(\phi_k)$	$P(X_{ki} = 1;\ \phi_k)$, the probability of a detect state, D_k, given a target in L_k and ϕ_k is allocated there.
$q(\phi_k)$	$P(X_{ko} = 0;\ \phi_k)$ the probability of a nondetect state, $\overline{D}_k$, given a nontarget in L_k and ϕ_k is allocated there.
CW	Correct whereabouts response.
CD	Correct detection response.
R_i	Choice of L_i in whereabouts task.
FA	Focused attention.
DA	Divided attention.

Representation Assumptions

1. $\Phi \geq$ Total capacity is positive.
2. $\Phi = \Sigma\phi_k$ Capacity is additive or conserved.
3. $b(\phi_k)$ is right continuous, concave and increasing.
4. $q(\phi_k)$ is right continuous, concave and increasing.

In our application of this model we have chosen an empirical situation for which it makes sense to assume that $b(\phi_k) = q(\phi_k)$ for all ϕ_k. In other applications it might make sense to assume that X_{ki} always starts out with the value zero and with some probability changes to the value 1. In other applications there might be reason to assume that $q(\phi_k) < b(\phi_k)$ or that $b(\phi_k) > q(\phi_k)$ for all ϕ_k. It is also possible that the relation between $b(\phi_k)$ and $q(\phi_k)$ changes with ϕ_k.

PROBABILITY OF A CORRECT WHEREABOUTS SEARCH

Here we treat only the case of two locations. In our experiment these are the positive and negative diagonals.

Focused Attention. This is also called the whereabouts strategy: attend to one diagonal and if fail to detect the target, then guess the unattended diagonal. It follows that

$$P(CW/FA) = P(L_k)q(\Phi) + P(L_j)b(\Phi)$$

where L_j is the attended diagonal and L_k is the ignored diagonal. For convenience we assume $P(L_1) \geq P(L_2)$. The two possible whereabouts strategies are obvious and, for each, the probability of a correct whereabouts response are given by

1. Attend to L_2:

$$P(CW/FA) = P(L_1)q(\Phi) + P(L_2)b(\Phi);$$

2. Attend to L_1:

$$P(CW/FA) = P(L_1)b(\Phi) + P(L_2)q(\Phi).$$

What is the better strategy? If either $P(L_1) = P(L_2)$ or $q(\Phi) = b(\Phi)$ it does not matter and as we show later the whereabouts (FA) strategy gives a higher P(CW) than the divided attention strategy. If we assume $P(L_1) \geq P(L_2)$, then which location should be ignored depends on the relation between $q(\Phi)$ and $b(\Phi)$.

Divided Attention. Here the strategy is to divide attention among the locations in order to maximize the probability of a correct detection response and then choose the location most likely to contain the target (i.e., has the highest posteriori probability for the target given the pattern of detect states and the amount of attention allocated to each location. Choosing the location with the greatest posterior probability is the best strategy given one must divide attention (i.e can't use the whereabouts strategy) and then try to maximize P(CW/DA); therefore, it is sufficient to show that the best P(CW/DA) is always smaller than the best P(CW/FA).

The probability of a correct whereabouts response given divided attention is

$$P(CW/DA) = P(R_1|L_1)P(L_1) + P(R_2|L_2)P(L_2)$$

Now, $P(R_1 | L_1)$ must be decomposed into four components:

$$P(R_1 | X_{11} = 0, X_{20} = 0)\ P(X_{11} = 0, X_{20} = 0|L_1),$$

$$P(R_1 | X_{11} = 1, X_{20} = 0)\ P(X_{11} = 1, X_{20} = 0|L_1),$$

$$P(R_1 | X_{11} = 1, X_{20} = 1)\ P(X_{11} = 1, X_{20} = 1|L_1),$$

$$P(R_1 | X_{11} = 0, X_{20} = 1)\ P(X_{11} = 0, X_{20} = 1|L_1).$$

The same applies, of course, to $P(R_2 | L_2)$. For the sake of simplifying notation we shorten $X_{ki} = 0$ to $\overline{D}_k$ and $X_{ki} = 1$ to D_k.

Given a pattern of detect and nondetect states over locations, the choice of a location for the whereabouts response should maximize the possibility of a correct choice. The overall probability of a correct whereabouts response will then be a mixture of strategies that depend upon these patterns of detect and nondetect states. Thus, we first consider these patterns of detect states and the corresponding optimal choice. In all the cases below it is assumed that the optimal allocation plan (ϕ) for detection is used and there is sufficient Φ that this plan results in some positive allocation to each location. If there is not sufficient Φ, then the best choice of a location given failure to detect in the high probability location, L_1, is still the high probability location, L_1. This follows because under the optimal detection plan Φ is allocated entirely to L_1 and this means

$$P(L_1)\,[1-b(\Phi)] > P(L_2).$$

In other words, the conditional posterior probability for L_1 is still higher than the prior probability for L_2. (See Stone, 1975).

We now consider one by one the optimal choice for each pattern of detect states.

CASE 1. $D_1\,\overline{D}_2$ $(X_{1i} = 1 \text{ and } X_{2j} = 0)$

(a) $$P(L_1|D_1\overline{D}_2) = \frac{b(\phi_1)q(\phi_2)P(L_1)}{P(D_1\overline{D}_2)}$$

(b) $$P(L_2|D_1D_2) = \frac{[1-q(\phi_1)][1-b(\phi_2)]P(L_2)}{P(D_1\overline{D}_2)}$$

The optimal strategy is to choose L_1 if $P(L_1|D_1\overline{D}_2) \geq P(L_2|D_1\overline{D}_2)$; otherwise, choose L_2. We consider the case where $P(L_1) = P(L_2)$. The optimal allocation of Φ is $\phi_1 = \phi_2$ so we have the decision rule choose L_1 if

$$b(\Phi/2)q(\Phi/2) \geq [1-b(\Phi/2)][1-q(\Phi/2)],$$

or when

$$b(\Phi/2) \geq 1-q(\Phi/2).$$

If $b(\phi) = q(\phi)$ then the decision rule reduces to choosing L_1 if

$$b(\Phi/2) \geq 1/2 .$$

CASE 2. $\overline{D}_1 D_2$ $(X_{1j} = 0 \text{ and } X_{2j} = 1)$

(a) $$P(L_1|D_1\overline{D}_2) = \frac{[1-b(\phi_1)][1-q(\phi_2)]P(L_1)}{P(D_1\overline{D}_2)}$$

(b) $$P(L_2|D_1D_2) = \frac{q(\phi_1)[b(\phi_2)]P(L_2)}{P(D_1\overline{D}_2)}$$

The optimal strategy is to choose L_2 if

$$q(\phi_1)b(\phi_2)P(L_2) \geq [1-b(\phi_1)][1-q(\phi_2)]P(L_1);$$

otherwise choose L_1. We consider the case where $P(L_1) = P(L_2)$ and so $\phi_1 = \phi_2$. The optimal choice of a location is L_2 if

$$q(\Phi/2)b(\Phi/2) \geq [1-b(\Phi/2)][1-q(\Phi/2)]$$

or when

$$b(\Phi/2) \geq 1-q(\phi/2),$$

If $b(\phi) = q(\phi)$ then the decision rule reduces to choosing L_1 if

$$b(\Phi/2) \geq 1/2.$$

CASE 3. D_1D_2 $(X_{1i} = 1 \text{ and } X_{2j} = 1)$

(a) $$P(L_1|D_1D_2 = \frac{b(\phi_1)[1-q(\phi_2)]P(L_1)}{P(D_1D_2)}$$

(b) $$P(L_2|D_1D_2) = \frac{[1-q(\phi_1)][b(\phi_2)]P(L_2)}{P(D_1D_2)}$$

The optimal strategy is to choose L_1 if

$$[b(\phi_1)][1-q(\phi_2)]P(L_1) \geq [1-q(\phi_1)][b(\phi_2)]P(L_2).$$

We first consider the case where $P(L_1) = P(L_2)$ so that $\phi_1 = \phi_2$. The optimal choice of a location is L_1 if

$$b(\Phi/2)[1-q(\Phi/2)] \geq [1-q(\Phi/2)]\, b(\Phi/2).$$

But these are equal so the choice is arbitrary.

CASE 4. $\overline{D}_1\ \overline{D}_2$ ($X_{1i} = 0$ and $X_{2j} = 0$)

(a) $$P(L_1 | \overline{D}_1\ \overline{D}_2) = \frac{[1-b(\phi_1)]q(\phi_2)P(L_1)}{P(\overline{D}_1\ \overline{D}_2)}$$

(b) $$P(L_2 | D_1\ D_2) = \frac{[q(\phi_1)][1-b(\phi_2)]P(L_2)}{P(\overline{D}_1\ \overline{D}_2)}$$

The optimal strategy is to choose L_2 if

$$P(L_2)q(\phi_1)[1-b(\phi_2)] \geq [1-b(\phi_1)]q(\phi_2)P(L_1).$$

Here we consider the case where $P(L_1) = P(L_2)$ and so $\phi_1 = \phi_2$. The optimal choice is L_2 if

$$q(\Phi/2)[1-b(\Phi/2)] \geq [1-b(\Phi/2)]q(\Phi/2).$$

But these are equal so the choice is arbitrary.

We now derive an expression for P(CW/DS) when $P(L_1) = P(L_2)$ and the optimal response selection rule is used for the case where $b(\phi) > 1-q(\phi)$.

$$
\begin{aligned}
P(CW/DS) &= P(L_1/L_1)P(L_1) + P(L_2/L_2)P(L_2) \\
&= P(D_1/L_1)P(\overline{D}_2/L_1)P(L_1) \\
&+ P(D_1/L_1)P(D_2/L_1)P(L_1) \\
&+ P(\overline{D}_1/L_2)P(D_2/L_2)P(L_2) \\
&+ P(\overline{D}_1/L_2)P(\overline{D}_2/L_2)P(L_2)
\end{aligned}
$$

Substituting the model expressions in the equation above, we have

$$
\begin{aligned}
P(CW/DS) &= b(\phi_1)q(\phi_2)P(L_1) \\
&+ b(\phi_1)[1-q(\phi_2)]P(L_1) \\
&+ q(\phi_1)b(\phi_2)P(L_2) \\
&+ q(\phi_1)[1-b(\phi_2)]P(L_2) \\
&= b(\Phi/2)P(L_1) + q(\Phi/2)P(L_2) \\
&= .5[b(\Phi/2) + q(\Phi/2)]
\end{aligned}
$$

Comparison of P(CW/DA) and P(CW/FA).

$P(CW/FA) = .5[b(\Phi) + q(\Phi)]$ and

$P(CW/DA) = .5[b(\Phi/2) + q(\Phi/2)]$

Since $b(\Phi) > b(\Phi/2)$ and $q\ (\Phi) > q(\Phi/2)$ it follows that

$.5[b(\Phi) + q(\Phi)] > .5[b(\Phi/2) + q(\Phi/2)]$ or

$P(CW/FA) > P(CW/DA)$

The reverse prediction obtains, however, if $b(\phi)<1-q(\phi)$: that is if conditional probability of a detect state given the target is present is smaller than the probability of this state when the target is not present. We leave this proof to the reader.

CONTINUOUS STATE MODEL

Additional Notation

Symbol	Description
$f_t(X)$, $f_d(X)$	The probability density function for target and distractor random variables, respectively.
μ_s, μ_n	The mean of the target and distractor random variable, respectively.

Representation Assumptions

1. $\phi \geq 0$	Total capacity is positive
2. $\phi = \Sigma\phi_k$	Capacity is additive (conserved under partitions of it)
3. $\sigma^2_{X_{k1}} = \frac{\sigma^2}{\phi_k}$	σ^2 is the variance inherent in the stimulus representation.
4. $f_d(X) = f_t(X+C)$	The probability distribution of X_{k1} is a shift of the probability distribution of X_{k0}.
5. f_d, f_t	are symmetric, unimodal distributions with finite means and variances.

For convenience we assume that $\mu_n = 0$. The assumptions that f_d, f_t are symmetric, unimodal with the same variance inherent in the stimulus representation are analogous to the assumption that $b(\phi_k) = q(\phi_k)$ in the two-state model.

PROBABILITY OF A CORRECT WHEREABOUTS SEARCH

We first treat the case of two locations and then the case of n locations.

TWO LOCATIONS

Focused Attention (whereabouts strategy). The subject is instructed to attend to one diagonal and if the target is not detected, then guess the unattended diagonal. For the continuous state model we assume that the subject picks a criterion β_k for the attended location L_k and if $X_{ki} < \beta_k$ then the unattended location is selected for the whereabouts response. Thus, if L_k is the attended location

$$P(CW/FA) = P(X_{k0} < \beta_k)[1-P(L_k)] + P(X_{k1} > \beta_k)P(L_k).$$

We denote by G(Z) the standardized cumulative distribution function and then write

$$P(CW/FA) = G\left[\frac{\beta_k - \mu_n}{\sigma/\sqrt{\Phi}}\right][1 - P(L_k)] + \left[1 - G\left[\frac{\beta_k - \mu_s}{\sigma/\sqrt{\Phi}}\right]\right]P(L_k)$$

It is assumed that $P(L_k) = 1-P(L_k)$ and β_k is optimal ($\beta_k = \frac{\mu_s}{2}$ will maximize the probability of a correct detection at the attended location), then it follows that

$$P(CW/FA) = G\left(\frac{-\sqrt{\Phi}\,\mu_s}{2\sigma}\right)$$

Divided Attention (detection strategy). The subject is instructed to attend equally to both locations (diagonals) and choose one location (diagonal) most likely to contain the target given the impressions from the brief flash. In the model we assume that the subject chooses the location with the largest random variable. Thus,

$$P(CW/DA) = P(L_1)P(X_{20} - X_{11} < 0) + P(L_2)P(X_{10} - X_{21} < 0).$$

Again, letting G denote the standardized distribution function, the equation above can be rewritten as follows:

$$P(CW/DA) = P(L_1)\ G\left(\frac{0 - \mu_s}{\sqrt{\sigma^2/\phi_1}}\right) + P(L_2)\ G\left(\frac{0 - \mu_s}{\sqrt{\sigma^2/\phi_2}}\right).$$

If it is assumed $\phi_1 = \phi_2$ because $P(L_1) = P(L_2)$, then we have

$$P(CW/DA) = G\left(\frac{-\sqrt{\phi}\mu_s}{2\sigma}\right).$$

Comparing this final form for P(CW/DA) and the final form of P(CW/FA) we see P(CW/DA) = P(CW/FA). If subjects use a nonoptimal β in the focused attention condition we will have P(CW/FA) < P(CW/DA).

N LOCATIONS

In the case of N locations we have Monte Carlo results for the Gaussian probability function and the assumption that $P(L_k) = P(L_j)$ for all k and j.

Whereabouts Strategy. Here one location L_k is chosen to be ignored; if for all $l \neq k$, $X_{li} < \beta_l$, then L_k is chosen for the whereabouts responses. If however, there are some X_{li} such that $X_{li} > \beta_l$, then L_r is chosen such that $L_r = \max \{X_{li}\}$. The probability of a correct choice under this strategy is

$$P(CW/WS(FA)) = P(L_k) \prod_{l \neq k} G\left(\frac{\beta_l}{\sigma/\sqrt{\phi_l}}\right)$$

$$+ \sum_{l \neq k} P(L_l) \int_{\beta_k}^{\infty} g\left(\frac{X - \mu_s}{\sigma/\sqrt{\phi_l}}\right) \prod_{k, i \neq l} G\left(\frac{X}{\sigma/\sqrt{\phi_i}}\right) d_x \, .$$

Detection Strategy. Again, it is assumed all locations are equally likely to contain the target. The location chosen, L_k, for the whereabouts response is assumed to be the one such that $X_{ki} = \max \{X_{lj}\}$. Unlike the whereabouts strategy, attention may be allocated to all locations so that the probability of a correct detection response is maximized. Since $P(L_k) = P(L_l)$ for all k,l this leads to $\phi_i = \phi_k$ for all l,k. We now write the probability of a correct whereabouts response as a function of this detection strategy (DS) as

$$P(CW/DS(FA)) = \int_{-\infty}^{\infty} g\left(\frac{X - \mu_s}{\sigma/\sqrt{\phi}}\right) G\left(\frac{X}{\sigma/\sqrt{\phi}}\right)^{n-1} d_x$$

In Table 4 we present some sample computations of P(CW/WS) and P(CW/DS) for various values of μ_s, σ, β_k, Φ, and n (number of locations). These computations show that dividing attention equally among the cells and then choosing the maximum random variable provides a higher probability of a correct whereabouts response than the whereabouts strategy in which a location is ignored.

FIGURE CAPTIONS

Figure 1. The relation between $\underline{P}_2$ (probability of a correct location judgment in the two-location task) and $\underline{P}_4$ (same probability in the four location task) predicted by several models and data obtained by Shaw (Note 1) in a visual search experiment.

Figure 2. Predictions of the Capacity Allocation Model: the difference between P(CW/FA) and P(CW/DA) as a function of the difference in the false alarm rate for focused versus divided attention for two different values of total capacity, Φ.

Figure 3. The relation between $\underline{P}_2$ (probability of a correct location judgment in the two-location task) and P_4 (same probability in the four-location task) predicted by the boundary curve and two attention models when the choice is between diagonals in the four-location task. The data points are from four subjects in the present experiment.

Figure 4. The relation between P(CW/DA) and P(CW/FA) for the Sample Size Model and the Capacity Allocation Model. In general, the class of continuous state models predicts data on or below the diagonal line, while the class of two-state models predicts data above the diagonal line. The data points are from the four subjects in the present experiment.

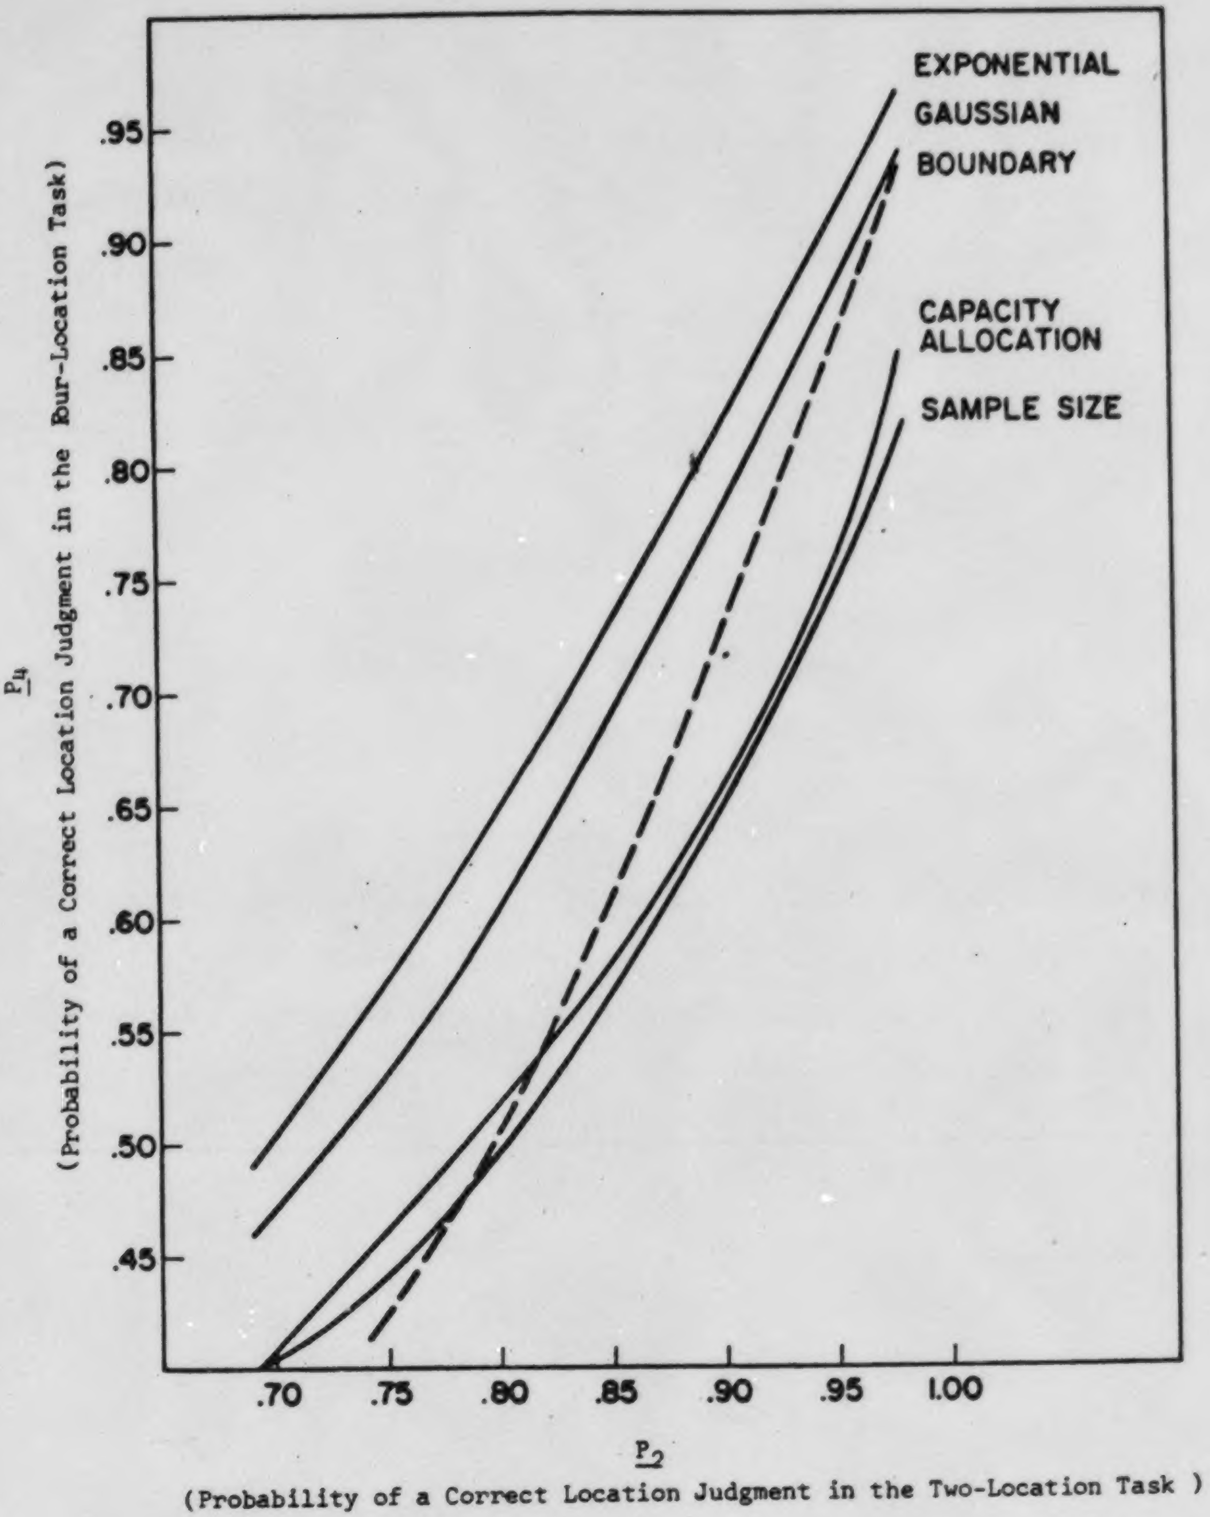
EXPONENTIAL
GAUSSIAN
BOUNDARY
CAPACITY ALLOCATION
SAMPLE SIZE
.95
.90
.85
.80
.75
.70
.65
.60
.55
.50
.45
P_4
(Probability of a Correct Location Judgment in the Four-Location Task)
.70
.75
.80
.85
.90
.95
1.00
P_2
(Probability of a Correct Location Judgment in the Two-Location Task)

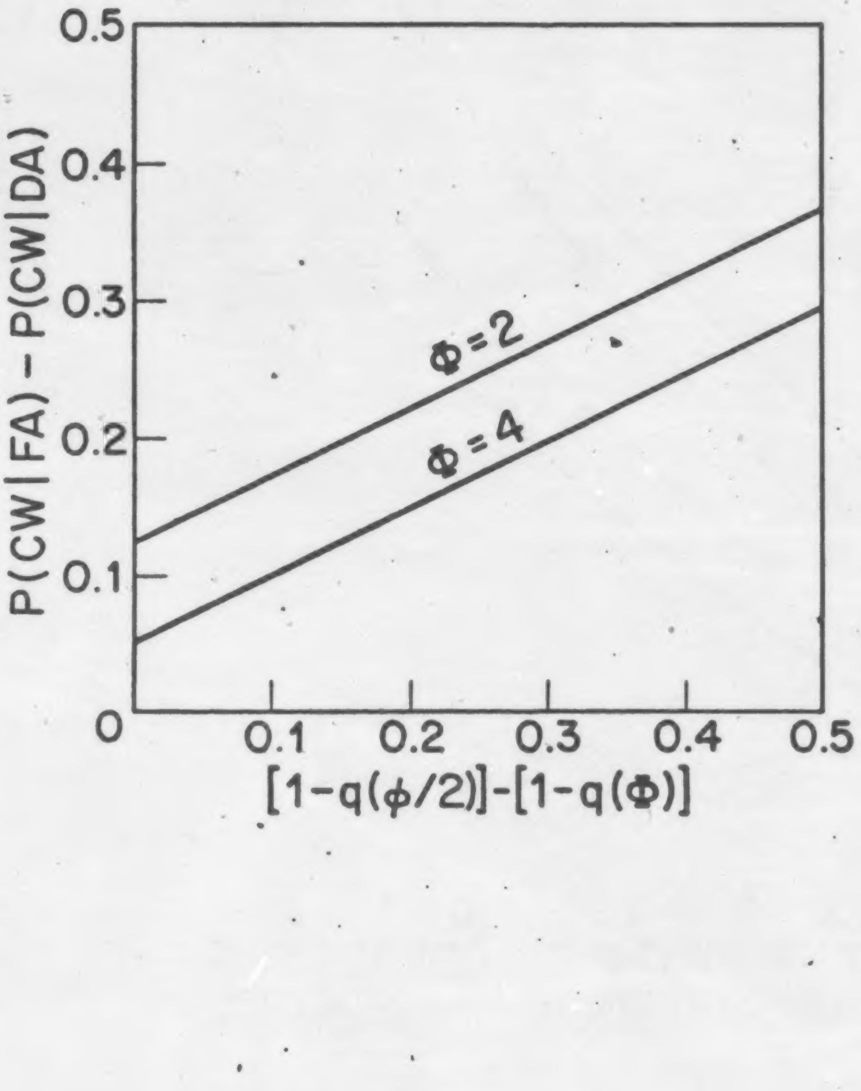

343

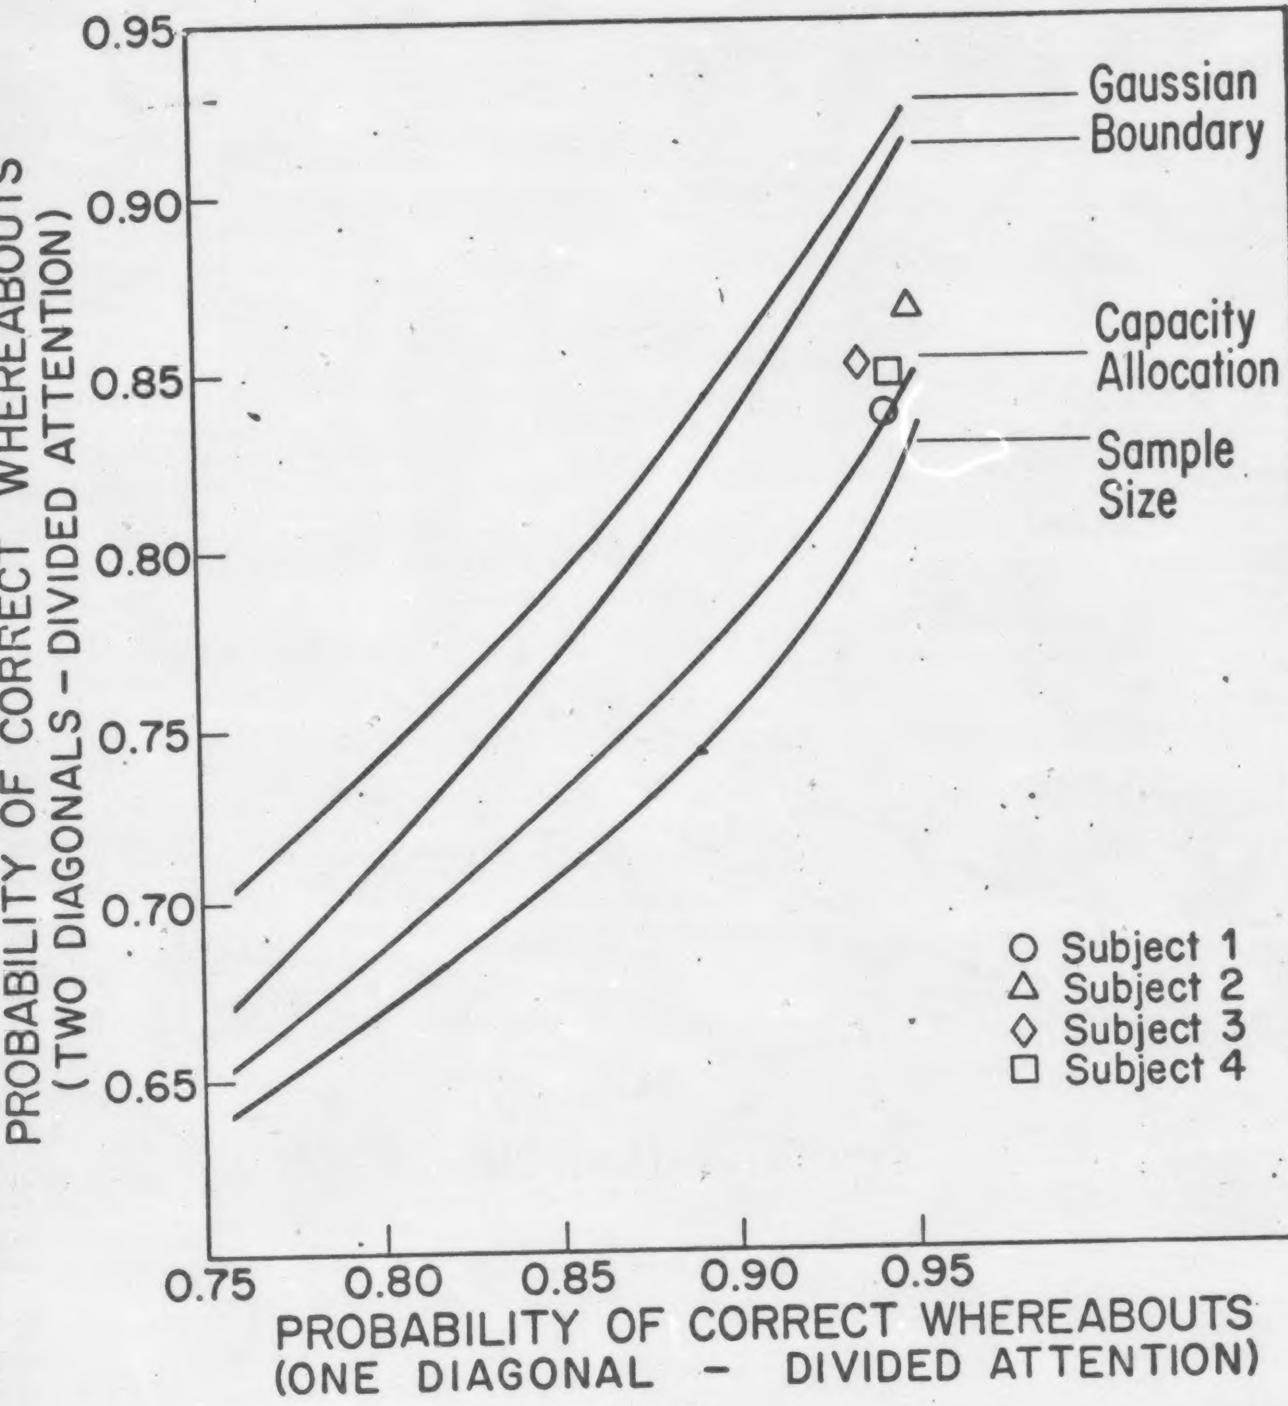

0.95
0.90
0.85
0.80
0.75
0.70
0.65
PROBABILITY OF CORRECT WHEREABOUTS (TWO DIAGONALS - DIVIDED ATTENTION)
Gaussian
Boundary
Capacity
Allocation
Sample
Size
○ Subject 1
△ Subject 2
◇ Subject 3
☐ Subject 4
0.75
0.80
0.85
0.90
0.95
PROBABILITY OF CORRECT WHEREABOUTS (ONE DIAGONAL - DIVIDED ATTENTION)

344

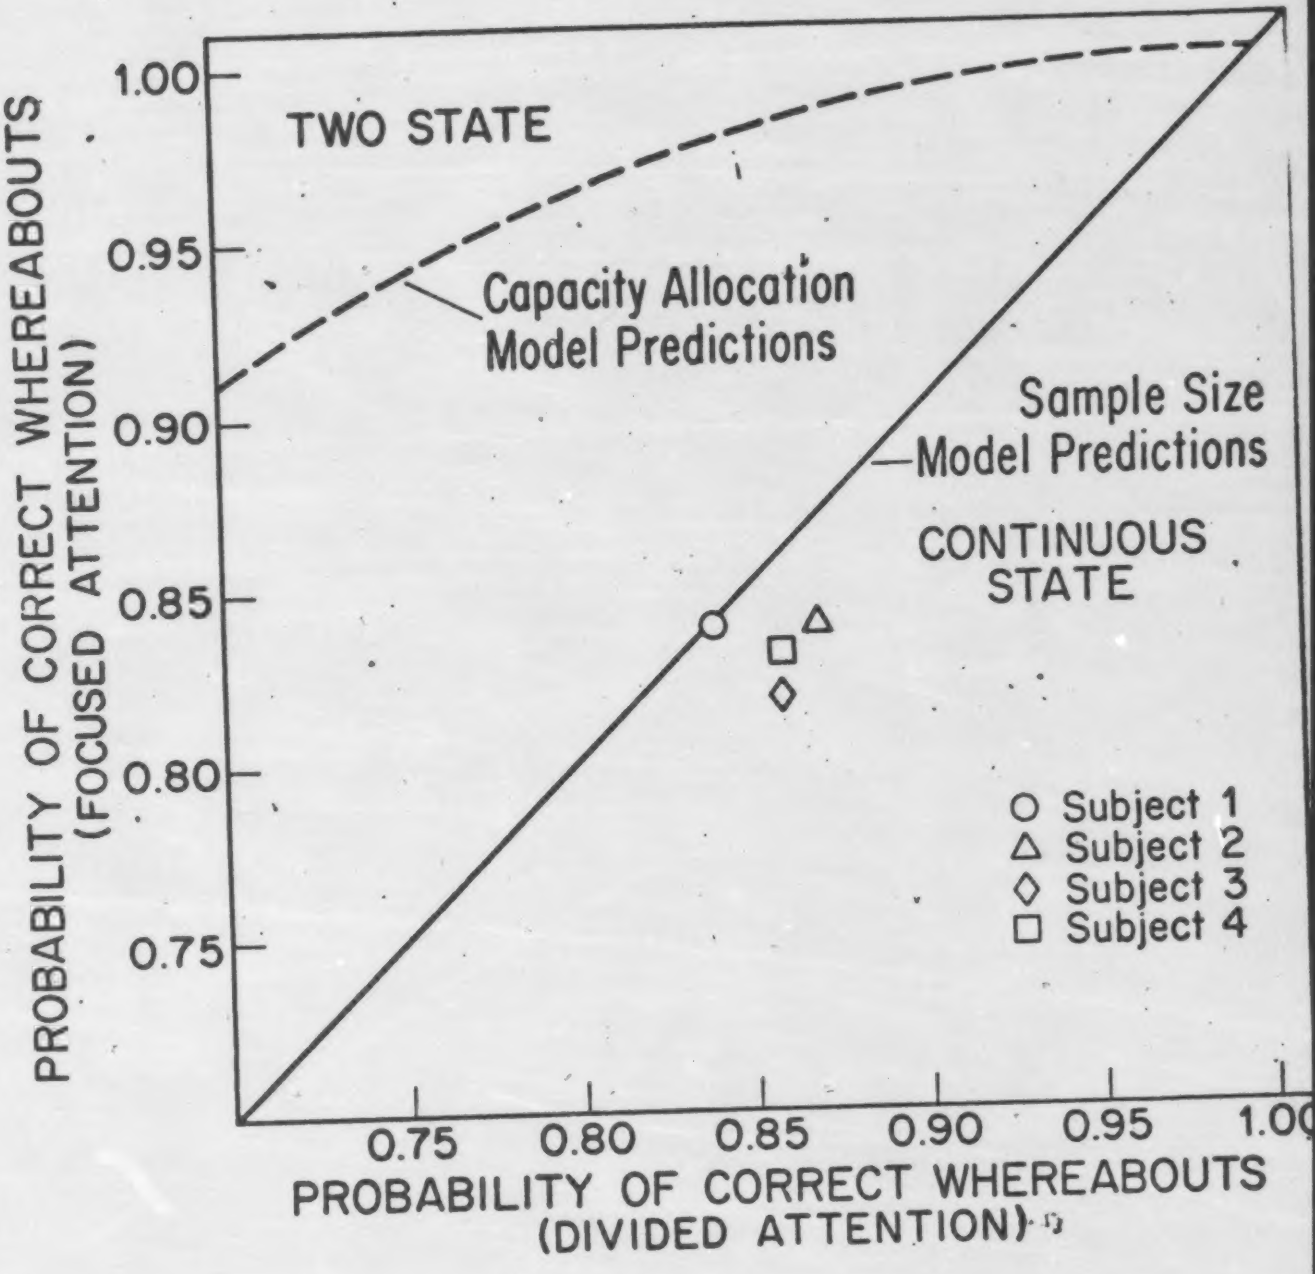
1.00
0.95
0.90
0.85
0.80
0.75
PROBABILITY OF CORRECT WHEREABOUTS
(FOCUSED ATTENTION)
TWO STATE
Capacity Allocation
Model Predictions
Sample Size
Model Predictions
CONTINUOUS
STATE
○ Subject 1
△ Subject 2
◇ Subject 3
□ Subject 4
0.75
0.80
0.85
0.90
0.95
1.00
PROBABILITY OF CORRECT WHEREABOUTS
(DIVIDED ATTENTION)

AUTOMATISM AND ATTENTIONAL PROCESSES

Lake Wilderness Attention Conference

Richard Shiffrin

I'm going to use the opportunity today to engage in some general speculations about the role of automatism and attention processes. I'm going to review some old work, discuss some ways in which the notions of automatism help explain traditional findings in attention, and present some new work.

To begin with, it's necessary to define the notions of automatism and controlled processing. This gets into some very difficult issues immediately, because I don't think anyone has yet managed to come up with necessary and sufficient conditions to define automatism or distinguish it from controlled processing (although everyone seems comfortable in using these terms). Among others here, I have proposed a dichotomous model in which automatism and controlled processing are qualitatively different processes. We put forward first a capacity definition: "any process that does not use general processing resources and does not decrease the general processing capacity available for other processes is automatic." It seems reasonable that any such process should be automatic. The second definition is also a useful condition: "any process that demands resources in response to external stimulus input, regardless of the subject's attempts to ignore the distraction, is automatic." This is a rather strong condition; one might wish to weaken it by adding an attentional condition that says that a process is automatic if it occurs when the subject is not attending to it (or trying to initiate it), as opposed to trying to stop it. That is, one can easily imagine processes that the subject can successfully stop but are automatic in the sense that they take place when the subject is not attending. A third possible criterion involves learning. Most automatic processes exhibit some sort of gradual and continued improvements in performance over long periods of practice, and the process of automatization might be defined by this characteristic.

There are numerous problems in trying to use any of these definitions to establish a difference between automatic and controlled processes in practice. Some of the difficulties arise because almost any task is carried out with a mixture of controlled and automatic processing, both in the sense that each type of process can initiate the other, and also in the sense that both can be carried out simultaneously. Furthermore, these difficulties are increased by the fact that the automatic component shows some sort of continuity as a function of practice during the gradual course of learning.

In spite of these theoretical and practical difficulties, which I don't want to go into today, it is often possible to determine, the automatic and controlled components in the context of a specific experimental paradigm. One such paradigm was used a few years back by Wally Schneider and myself. It is probably worth mentioning very briefly a couple of those results. The basic idea is that a search task is carried out by the subject. Some memory set items to be searched for are presented at the beginning of a trial, then a series of characters appear in a frame. The subject says whether any of these memory characters appear in the frame, yes or no. There's two main conditions of interest, varied mappings (VM) and consistent mappings (CM). In the varied mapping condition, over trials the targets and distractors trade roles from trial to trial; targets in one trial are distractors in another and vice versa. In this condition, very little can be learned. In the consistent mapping conditions, the targets and distractors remain fixed over long periods of trials; the targets are always targets, the distractors are always distractors. Following are the typical results. These are the reaction time results when the subject is presented a single frame and is asked to respond as quickly as possible, yes or no. The dashed lines are the consistent mapping conditions in which you see that there's very little effect of load. However many items the subject is searching for, or however many items there are in the display affect the reaction time very little. The subject is clearly doing something very parallel and independent of the load in the CM condition. However, the varied mapping condition gives rise to the traditional search results: a more or less linear increase in reaction time as a function of load.

Very similar results occur in the case of the multiple frame procedure in which accuracy is the measure rather than reaction time. In the multiple frame procedure, something like 20 frames in a row are presented rapidly and the subject is asked to say whether there is a target anywhere in that series of frames. Now the question is: does the subject have time to find the target (rather than how long does it take to find the target). Thus we examine the probability of finding the target (instead of the reaction time). Otherwise the situation is the same, and very similar results come out. I won't go over these in detail, but the basic idea is that in the varied mapping condition, the frame time which allows a high probability correct is very large. 800 milliseconds per each frame is needed if there's a high load, (16) to produce 70% correct. On the other hand, in the consistent mapping conditions, these times per frame are all very fast, and load doesn't matter very much.

We were able to fit the varied mapping results in both the multiple frame and single frame conditions reasonably well with the serial terminating comparison process model similar to Sternberg's except that termination is assumed rather than exhaustion of the search. What I'm going to say a little more about today, is what's happening in the consistent mapping conditions where we assume some sort of automatic detection is occurring. The basic idea that Schneider and I proposed was that some sort of attention process was occurring in the consistent mapping condition such that the target items, were attracting attention to themselves, and thereby bypassing the necessity for a serial search.

One source of evidence for this model arises from an attention study. The subject is given a relevant diagonal on which he carries out a varied mapping search, let's say for letters in letters (letter targets and letter backgrounds); there's an irrelevant diagonal in which letters are appearing but they're to be ignored all the time. As a series of multiple frames are presented, let's say the relevant target is "F". Suppose "3" is an old learned consistent mapping target that has been trained previously. Suppose the "3" occurs in the same frame as the target that's to be detected, the "F". We compare this case with the case when the old trained target, the "3", occurs on the irrelevant diagonal in the frame prior to, or the frame subsequent to, the actual target to be detected. Then the question is, "Will the presence of this previously trained item, even if it's in a "to be ignored" location, hinder the detection of the target on the relevant diagonal, even though the subject knows what to attend to and is supposed to ignore the other diagonal?" The results are as follows: When the relevant target is in the same frame as automatic item, it harms detection considerably; it drops detection accuracy by about 20% (which doesn't happen when the CM item is before or after the relevant target). The explanation holds that the CM target is calling attention to itself, so that the attention system orients away from the relevant target which is missed as a consequence.

349

Question: Does your subject know in advance, he might get a situation like this?

Answer: Yes, these numbers are occurring somewhere on every trial, in some frame. They're always to be ignored, so the subject knows that they're occurring, and he knows to ignore them; it's not a surprise. They're occurring regularly.

Sometimes they're in the same frame as the relevant target, sometimes they are somewhere else entirely. But he knows they're there, and he knows he's supposed to ignore them.

Question: Is the attention attracting capability lost as your experiment progresses?

Answer: Not over the course of the session or two over which we ran this experiment. Of course, if we ran this study long enough, the subjects would adapt eventually. That is, the tendency for this item to attract attention would undoubtedly eventually go away, but we didn't run long enough to have that.

There are some other relevant results concerning the learning of an attention attracting response. One somewhat peripheral finding, but worth pointing out, is a study that Wally Schneider carried out. The results I just showed you indicated a deficit in detection when an automatic target occurred at the same time as a relevant VM target. Schneider's results occur when the subject is asked to carry out two tasks at the same time, one VM and one CM. One search task is automatic (on one diagonal) and on the other diagonal the subject is asked to carry out a VM search. But the two targets do not occur at the same time. The subject is simply required on every trial to find any target that occurs (either might occur but only one does). The results show that the subject is able to carry out both search tasks as well as when control conditions are used in which only one task is required. This shows that the subject is able to carry out both tasks together, as well as either alone. This result only occurs when the subject is told to focus all his attention on the VM task. If the subject is not told what to do, or is asked to divide his attention between the two diagonals, he does as well on the CM task as before (although his response bias changes) but because he's withdrawing attention from the controlled, the VM task suffers greatly.

This result shows you can carry out VM and CM tasks at the same time. It indicates that at least one of these tasks is not demanding capacity, so we can do the other one. Clearly, the one that's not demanding capacity is the automatic task, because that's the one that you don't have to pay any attention to by instruction to get these results. So, apparently the automatic detection does not require an attentional effort.

Question: They don't need to attend to the automatic diagonal and it harms them if they do so?

Answer: Yes. That's absolutely right. They don't know that they don't need to devote attention to the automatic diagonal. Hence, if you leave them alone, they will devote unnecessary attention to the automatic task, attention that is not really needed. It may even be harmful to the CM task to do so, but that needs to be explored. What is clear is that the VM task is harmed when (unnecessary) attention is devoted to the CM diagonal.

Question: Have you got those results?

Answer: Oh yes, I've got them in my briefcase; I can show them to you later. There's clearly a big deficit when they're competing with each other.

There are some reversal results from the '77 paper which are worth mentioning briefly for their attentional implications. The subject starts off learning one set of letters as targets in another set of letters as distractors in a CM procedure. Performance gradually improves in a multiple frame dectection task until asymptotic performance is reached. At that point the targets and distractors are reversed. That is, if the subject has been learning A targets in B distractors, he's now told to search for B targets in A distractors. The two sets are simply reversed. What happens is a severe decrement. The subject not only is not transferring, but has negative transfer. The subject drops below the original level of performance seen at the start of training. So it's actually harmful to reverse after learning. One might imagine that if the targets are attracting attention, that they will continue to do so after the reversal and pull attention away from the new targets, which are the old distractors. It's important to realize that in this study, there were only two characters in each frame. Thus after reversal, there was only one new target and one "old" target in the target frame. In many of the studies where we get this negative effect of reversal, we find that occurs when in the target frame there's just one of these old trained targets. If there's many old trained targets in the target frame, we don't get this result. That is, performance just returns to the VM level after reversal. Perhaps when there are many targets all attracting attention simultaneously, attention doesn't go anywhere. At least that's a possibility. But in this case, we do get a negative effect, again interpretable in terms of attracting attention.

One last result, again with respect to this whole issue of attention attraction, is a study Sue Dumais carried out on transfer of training. She used a transfer of training design that is basically fairly simple. The subject is given CM training on A targets with B distractors and C targets with D distractors. The subject learns these well. At the same time they're trained on VM tasks using other characters. After transfer, the targets can be kept the same by keeping the old CM targets and using VM background items. Alternatively, the old background items can be kept the same, and new VM items can be used as targets.

The basic question is: What is being learned in original CM training? The targets? The distractors? Both?

A reaction time task was used. These bars show the difference in reaction time between frame sizes of 16 and frame sizes of 4. In effect, these are the slopes of the search functions. A large slope occurred in initial VM framing, and a small slope occurred in original CM training. The results after training can be summarized by saying that both target and distractor transfer was excellent. That is, whether the targets were kept the same, or whether the background items were kept the same, very good transfer occurred, about as good as a continuation of CM training.

Question: Why are the CM slopes not zero?

I'm going to come back to this and say a few words about it later. We find, in much of Sue Dumais' work, that CM subjects show a slope that wasn't present in the original Shiffrin and Schneider results. The subjects that Schneider and I ran originally, did not use what I'd call a great deal of effort in the CM conditions. If you're thinking of effort in Kahneman's old terms, let's say.

That is, they ran in pairs and they talked with each other, discussed the events of the day, world affairs, their homework, and other such things while they were carrying out the automatic detection. They showed this flat reaction time curve. Sue Dumais' subjects seemed to be a little bit more motivated. They were trying fairly hard in the CM conditions, even though it was an easy task for them. We think that, although it wasn't needed, they were carrying out controlled search in addition to the automatic detection that was all they actually needed. As a result, by mixing these two kinds of search, they produced non-zero slopes. Perhaps on some trials they were using controlled search, and on some trials they were using the automatic detection. By mixing the two, you get a small slope with a two to one slope ratio, as opposed to a flat curve. We think that's what's happening.

Question: Were they as well practiced as in the earlier studies?

Answer: Yeah. They were practiced quite well, certainly to asymptote.

Question: Why does mixing produce linear RT functions?

Answer: Well, it depends how you mix. There's a lot of different ways to do this. A number of different models can be proposed, and this gets a little technical. If you mix on a random proportion of trials then these linear results will be predicted. Other kinds of mixing are a little more problematical. For example, if they're competing in parallel, and the fastest wins, things get complicated, and then it gets much harder to find a version that will predict these results. There's a number of different versions of models that have to be considered in that respect. There are some models of a more complicated kind that still handle these results, but it's much more difficult to find them. The easiest thing to do is to imagine that the subjects, for one reason or another, are randomly using one or the other, on some proportion of trials. Perhaps they try both, but some random event determines which gets used. Whether that's a good assumption or not, I don't know. We're carrying out, incidently, a test to find out if it indeed is the case that subjects are mixing controlled search into the CM conditions. We're carrying out an experiment in which we give them a subsidiary task to do that will take their attention and cause them to devote it to something else, thereby removing it from this task and hopefully removing this slope. We don't know the results yet.

How do we explain the transfer results? Well, the model we have in mind is the same sort of model I've been proposing in which attention attraction is learned to these different items. The notion is that during training attention responses will become attached to items. The CM targets get very strong tendencies to attract attention, and the distractors become very weak in their tendency to attract attention, relative to the neutral amounts of tendency to attract attention that VM items get. Then in transfer you can explain the good performance for either background transfer or target transfer by the fact that one of these items will have a stronger tendency to attract the attention than the background. Whether you keep the background the same and put new VM items in as targets, or keep the targets the same, and new background items in as distractors, one of these will stand out with respect to the others because it will have a stronger tendency to attract attention. That explains the good transfer results. You might ask, "Do we know that VM items have an intermediate tendency to attract attention?" The answer is yes. Sue Dumais carried out another series of conditions which show the same transfer results if, instead of VM items in these transfer tasks, totally novel items are used. That is, new items the subjects haven't seen before are equivalent to VM items given much training. That indicates that the amount of tendency to attract attention that the VM items have is about the same as novel items have. That is, the result of VM training is to leave the tendency to attract attention at a constant level.

In summary, this whole picture looks fairly simple and fairly consistent. The items that are trained in a consistent manner as targets are developing a tendency to attract attention. The items that are trained as distractors are going the opposite way: they're developing a tendency not to attract attention. This picture looks fairly straightforward, and fairly simple. It doesn't explain everything, but as a summary result, it presents a fairly simple and accurate picture.

Let me just say a few words about the implications of these kinds of results for general attention models and processes. We would want to argue that the bottleneck that's seen in a large class of selective attention studies, that Broadbent originally described with a filter model, is identified with the limited serial comparison search process that we've been talking about in these studies. Subjects go through these items one at a time, comparing them, the limited, serial-comparison, search process may be identified with the bottleneck in the filter model. On the other hand, we suggest that all the stimuli are processed automatically to very high levels, up to perhaps meaning in some cases, so that processing doesn't stop at a low level. In that sense, the model is very much like that of say, Deutsch and Deutsch. So this is really a hybrid approach, a combination of the filter model, if you like, and the Deutsch & Deutsch approach. Finally, consider the kinds of evidence that led Anne Treisman to propose her modification of the filter model, the attenuation approach. Her evidence mostly consisted of demonstrations that the stimuli in not-to-be-attended locations, occasionally receive processing or cause effects. To a large degree, the kinds of results that I've been presenting here suggest that such results can be understood as the result of the "to be ignored" stimuli automatically, attracting attention for one reason or another. Such stimuli might have a learned tendency to attract attention (such as a person's own name in a to-be-ignored location) or might have an innate tendency to attract attention (such as an unusual color or loud noise causing a startle response). Many of the results in which to-be-ignored stimuli are nevertheless given processing may be due to their tendency to attract attention automatically.

This theoretical approach uses automatism as a basis for understanding selective attention. The approach has the merit of combining the research in selective attention with that in the search domain, and can explain both at the same time.

Most of the recent studies we've carried out have used above threshold stimuli, and one might ask "what happens if you try to extend this reasoning and this model to near threshold stimuli?" The reason this is an important question is that one might suspect automatism or automatic detection doesn't work very well for stimuli which are perceived so poorly that they are ambiguous. A number of years back, before I started the automatic work, I was involved with a series of threshold studies using what we called the successive-simultaneous procedure. Let me give just one typical example of a paradigm in this successive-simultaneous procedure. In the simultaneous condition, all the stimuli were presented simultaneously, preceded and followed by masks. In the successive condition, the same stimuli appeared one at a time, each preceded and followed by masks, and each presented for the same as they were in the simultaneous condition. (In some cases, we presented one diagonal at a time.) Subjects were asked to detect whether an "F" or "T" was in the display. It can then be asked whether detection accuracy is better in the successive conditions. If dividing attention among simultaneous positions causes fewer features to be extracted from a given character in a given position, then the simultaneous condition should result in poorer performance. Well, we carried out a dozen experiments along these lines, all of which showed that the simultaneous condition was essentially as good as the successive condition (or could be if you ran the experiments the right way) we therefore concluded that the amount of information extracted from a given display location was essentially equivalent in the simultaneous and successive conditions. This was evidence in favor of some sort of late selection model, rather than early selection model, because the information extraction was apparently equivalent in these two conditions.

But these results still raise a rather large puzzle, which is, "how could the simultaneous and successive conditions give equal performance?" Even if the amount of information extracted from each position is the same, somewhere along the line, attention is going to be operating. Even in short term store, the subjects have to scan these stimuli, and to decide about them; since the scanning and deciding is known to be capacity limited, one would still expect subjects to do better in the sequential condition. We once suggested that short-term store capacity wasn't exceeded, so that all the deciding and all the comparing that was needed could be carried out even in the simultaneous condition. That wasn't a very satisfying explanation.

Well, these automatic and controlled results provide one possible mechanism to explain the equality of performance. The idea is very simple. These simultaneous-successive studies were all run in CM conditions. The targets were always targets and they remained so. The idea then is that the subject developed attention responses to these CM targets. The difficulty of that explanation is that these stimuli are at threshold. That is, these stimuli are seen ambiguously sometimes, and how can this fact give rise to automatic detection? I now propose one explanation. The idea is simple. Sometimes under these threshold conditions, a letter is seen clearly; in those cases, the normal automatic attention response occurs, as one would expect in above threshold conditions. Sometimes one doesn't see the whole letter, but sees a unique pattern of features, a pattern that still serves to uniquely identify the letter. That unique pattern, during training, would also develop an automatic attention response, just as the name or the letter does, as long as it's unique. Thus unique partial features also call attention to the target location. The only really ambiguous cases occur when a feature or features are abstracted that do not serve to segregate targets from distractors consistently. In those cases, the automatic attention system can't operate, but on the other hand, in such cases using controlled processing won't help either. That is, if the target is ambiguous and distractors are ambiguous, it's not going to help to use control processing to search the display anyway. Thus automatic attention attraction can explain why the subject is able to exhibit equality of performance in the simultaneous and successive conditions. I should point out, it's very easy to test this argument. In fact, Dave Fogle has carried out a study at Indiana verifying this conjecture. He has carried out a simultaneous-successive study using both VM and CM training. In CM training, he found the usual equality ofsimultaneous and sequential conditions, but in VM training, he found a considerable advantage for successive conditions.

One supplementary point should be mentioned. John Duncan's results, that he reported and reviewed recently in Psychological Review, showed a deficit even in CM conditions for detecting simultaneously two targets occurring in the same frame. In those cases where he has two targets presented, he finds a deficit in the simultaneous condition relative to the sequential condition. That's explainable in our attention terms, as follows: when two targets are both calling attention to themselves, obviously both can't receive processing "first". At least one of them has to be delayed, in which case you're back into a situation where the successive condition should do better. Such reasoning can explain Duncan's results.

This argument, if accepted, does not affect the original conclusion from these studies, that the information extracted from a given location is the same in the simultaneous and successive conditions. That is, even if attention is being directed automatically to a given location, it's necessary that the goodness of the features extracted in the simultaneous and successive conditions be the same. If not, the attention attraction would be superior in the successive condition, and performance would be inferior in the successive condition. Thus the original conclusion we drew is still justified, but we are now interpreting the equality of performance in terms of a directing of attention due to an automatic response.

There's a number of other things that occur automatically in the processing system that affect attention and selection that perhaps ought to be touched on. The participants in this conference are going to be mentioning various of these. I want to mention just one in particular which might be called "automatic segregation" (or partitioning or grouping); this used to be studied often by Gestalt psychologists. In recent research, a series of studies by Anne Treisman has looked at this problem in the attention domain. I am raising the issue of some sort of automatic partitioning of a visual display into regions on one basis or another. This notion is not necessarily the same as attention being directed automatically to given regions, but simply an automatic segregation or partitioning.

As an example, consider these three displays. What we see is probably an automatic tendency for you to group these letters into two regions over here in the left panel, a series of columns over there in the middle panel, and a single zero in a surround of "X's" in the right hand panel. Note that there are really two things that have to be separated here. One is some tendency for automatic grouping to occur, and the second is perhaps some tendency for attention to be attracted to one of the grouped regions. Perhaps attention is attracted to this column, in this case, or perhaps attention is attracted automatically to that single zero, and perhaps there's some general rule that says attention tends to be attracted toward the smaller region rather than the larger region. But these are two separate issues; one is the tendency to automatically segregate or group the display, and the other is the tendency for attention to be directed toward one of these grouped regions. These are separable issues and can be experimentally studied. For example, if attention is being attracted to this isolated zero, one can check that by seeing if the presence of that zero harms search among all these "x's" and if there's a distance affect such that the closer a target is among these "x's" to the zero, the more harm to the search would take place. This is a very easy experiment to carry out. For example, if I asked you to find a "y", or tell me whether or not a "y" is somewhere in the "x's" as in this case, one could test whether the search time would be affected by the proximity of the "y" to the "zero". I think it's important that such tests be done, although I think that not too much has been done yet in terms of trying to separate those two issues.

Question: Are you saying that partitioning and grouping attract attention?

Answer: What I'm saying is that there are two separate processes: attention and grouping. Both may occur automatically, but they may not be completely identical. There may be a tendency sometimes for attention to be directed toward a region that has been automatically grouped. But these are two separate issues, and they should be experimentally separated so that you can distinguish the two.

Question: Does training determine the grouping?

Answer: Not in these examples. I think it's the nature of the stimuli that determine the grouping. This works best, of course, when you've got color displays, when colors are the basis for segregation. Anne Treisman has a number of results on that, having to do with color being very good. Shape is perhaps not quite as good, and you can go down the line. Regularities and irregularities in displays, oddities, have a good effect. But the main point I'm trying to make is that we should try to separate out the two aspects, the automatic attention attraction and the automatic grouping or segregation, both of which can affect attention and selection. One experiment that Sue Dumais carried out not too long ago, is worth pointing out, which gets at both these issues, and is worth mentioning.

What we did was train up "A" targets and "B" distractors in a consistent mapping condition. After training, we presented a field which contained both targets and distractors. Is there a tendency as a result of the training to partition, segregate or group this field into two regions or two sections, one containing the targets, and one containing the distractors? Sue Dumais carried out a counting test to get at this. We asked the subjects to count the the number of designated elements in the display, ignoring the background elements. In control conditions we didn't give consistent training; we simply asked the subject to count various things, like count the "V"s in a background of upside down "V"s.

That was a hard test, that was always very difficult. It took a long time per element. It took just as long to count numbers in letters or vice versa. However, it was easy to count O's and X's. That was very easy. 0's in x's is one of those things you can do automatically, somehow or another, and shows no slope effect of the irrelevant items. That is the number of the targets to be counted is a determinant of the reaction time, but the number of distracting elements has no effect.

Then, what we do is carry out CM training of letter targets in number distractors. After consistent training, we do the same experiment, and ask the same questions.

Let me show you some typical results here. These are abstracted from all the results. These conditions are a selection of those trials on which only one target was to be counted, so the answer was always "one". We have graphed reaction time as a function of the number of irrelevant items in the displays. After training of letters in numbers, using the CM procedure, we see these results: it is quite a bit easier to count letters in numbers and only a little bit harder to count numbers in letters. That is even numbers in letters is better than performance before CM training on letters in numbers.

We think that there's two effects going on. The reason both letters and number targets are helped is that there's some tendency for grouping of the display into two regions, targets versus distractors, which helps counting. In addition, there's this automatic tendency to attract attention, that causes you to do better on the trained targets, the letters, than the distractors, the numbers. We suggested that the result of training is that when you show a display of numbers and letters to a subject and ask him to count, his percept after training will look something like this in exaggerated form, with large letters standing out in a background of small numbers. This shows an accentuation of the targets and a diminution of the distractors. And we argue, it's easier to count the letters than it is to count the numbers, although it's the segregation that let's you count either set better than an homogenous set.

There's a variety of mechanisms by which this counting could take place. Note that the reaction times are very much slower than search reaction times in this study. These are like 150 to 200 milliseconds per count, so that they're doing something a lot slower than the normal serial scan (40 msec per comparison). They may even have time to scan over the whole display first to find all the targets, and then go back and count them. We don't know the exact mechanism, nor do we know for sure that the segregation that's occurring here is an automatic segregation or grouping, or whether it's a controlled segregation or grouping. There's enough time for the subject to carry out a search or a scan to segregate these displays into two halves, using some sort of controlled process if they want to. This is a matter for future research. Dan, do you have a question?

Question: Is this a matter of foreground-background perception?

Answer: Not exactly. We are using two types of characters so that in some sense there's a foreground and background, or something like that. I don't know if you want to call it that, it may be a foreground and background segregation. However, if we used three or four character types, then grouping may be into four regions with no one standing out especially, or with all of the CM trained characters standing out.

Question: Is training sufficient to produce this grouping effect?

Answer: That's a good question: some things are very difficult to train. Anne Triesman has shown that there's some things that are difficult to train, like conjunction of shape and color. Of course we'd have to train a long, long time, to be sure of what the limits are, but there's some controversy. Wally Schneider has some results indicating that for example, conjunction of features may be partially trainable, and Anne has some results that says, that she doesn't think so. At the moment it's a controversial issue, and perhaps I don't want to go into that today. What can be trained and what can't be trained is an interesting question, and I don't have too much to say about that right now. I think a lot of things can be trained, probably a lot more than we suspect, but there probably are limits, and I don't know yet what they are.

Question: Are your present results due to the use of letters and numbers?

Answer: Letters and numbers were used in this case, but none of our results in search processing lead us to think that there's anything special about letters and numbers in this respect, except that they're a little faster to train. We've always been able to duplicate all our results with letters and numbers, using letters and letters, and using highly confusing sets of stimuli, it just takes longer to train them. We used letters and numbers in this case because we didn't want to spend a few months of training, but none of our previous results leads us to think that there is anything special about these sets. It would be useful to carry out CM training, let's say on letters and letters, or some confusable sets, to see what the limits would be.

Question: Are the presented displays data limited?

Answer: That depends on what you mean by data limits. Our displays are not data limited in the usual sense since everything is seen clearly. The subject can clearly see all the characters. These are quite visible, quite contrasty, quite clear.

As a final conclusion, I want to say on the one hand that the notion of automatism and controlled processing is a fairly useful and valid dichotomy, and on the other hand that our theoretical dichotomy goes far towards explaining many of the mechanisms of selective attention, although obviously we're just beginning work in this whole area.

Question: Is there a special role played by color?

Answer: In part, that's answerable by experimentation, some of which has been done. I think that it undoubtedly is the case that color serves as a basis for discrimination, because it has some automatic properties like segregation, or attention attraction, or both, that will then carry selection effects with it. I don't think there's much question about that. I think that can be carried by innate properties like color discrimination or by trained properties. Maybe you could train any stimulus with enough training so it acts in the same way.

Question: Many studies have shown color to be a "special" cue of special salience, different from "response" sets.

Answer: Well, now, wait a minute Dan. This is an open experimental question. There is a general tendency for oddity to be able to serve as a cue in general.

If you only have two stimuli, red and blue or blue and red, and you switch back and forth, you might be able to use automatic segregation. However, it doesn't even have to be reds and blues, say it could be x's and o's or o's and x's, you can switch back and forth in the VM procedure, and might still get automatic segregation.

Question: What do you mean by oddity?

Answer: I mean some kind of perceptual oddity based on the stimulus pattern. There's clearly an issue of what happens with a small set of objects in a VM procedure. Something is very special about that, and I think that we can redefine the stimuli in those cases a little bit if we want, and still capture the attention attraction property, so we won't have to worry about the VM procedure. However, this business of color as being a special case is not at all clear. We're trying out an experiment with seven colors in the VM procedure. We rather suspect that when you increase the number of colors, we'll begin to get results that look more like the letter-number results. That is, in the VM procedure with seven colors, where you're told in a given trial which color to look for, we suspect the results will not look like the traditional color result that you can select out the one color you want from the whole set, and show no slope effect. With seven colors, it might be necessary to change for CM to VM to get the two different results. That is, there's something special I think about small sets that's not unique to colors, that may be unique to some kind of oddity result. This is an issue that's a little premature to talk about, because we're exploring this now. All I want to say is that the result that you're claiming is not at all clear, and it's an experimental question and we're looking at it now.

367

Question: Would zero slopes occur with any two different characters or dimensions?

Answer: We don't know. There's two issues here. One is the oddity question, and one is the VM or CM mapping, and both have to be explored.

Question: Can subjects learn to reverse quickly after much training?

Answer: I reserve judgment. We tried that. We took subjects and tried reversal, then reversed them back, reversed them back, and kept this up for a while. I think we've got it to about four once. We never carried it far enough to really be sure what was happening. It was clearly taking some time to reverse. Whether it was taking less time each time you reversed relative to some appropriate control and so on, well we never carried out an experiment with appropriate controls or carried it out far enough to be sure. It's a good question. Wally Schneider is studying this now.

Question: I thought automatism and controlled processing are dichotomous. How then can gradual learning occur?

Answer: There's two separate processes. One of them at least, the automatic one, shows continuous change with practice.

The point is, there's two separate tracks, there's two separate processes. One operates, perhaps, by attention attraction, and one operates by the subject volitionally assigning attention to different items in some sort of serial search. Thus there are two separate tracks that the subject can use to perform in these sorts of situations. The attention attraction may gradually develop over trials as you practice the subject, and eventually the subject will switch over to using it almost entirely. But, the fact that there's two separate processes does not mean that one can't develop with practice.

Question: What happens when the subject switches from controlled to automatic processing? How does the switch occur?

Answer: Well, there's basically two models. One model says that the controlled search being carried out in CM training in early trials is more or less irrelevant to the eventual production of the automatic response. In this view, VM search being carried out is only sort of a holding pattern that simply maintains the conditions necessary for the eventual development of your automatic response. In this view, controlled search simply allows the necessary conditions to take place, such that automatism will be developed. (We have carried out some experiments along these lines, and Schneider's carried out a few. We have studied things like how many targets you need to present versus how many distractors; how many times an item can appear as a target versus a distractor and still let the automatic attention response develop; how many times you need search for an item and find it versus not finding it. Things like this. We're trying to find out how fast automatism develops in these cases. That's another talk I give. I don't really have time to go into this now. It's another issue that would take another hour, so I don't think it would be wise to get into it now.) The other view is that somehow, something the subject is doing in the controlled search is a necessary condition for the eventual development of automatism. What that is we're not sure. In either case, we're not sure yet.

369

Question: Can the continuous model be distinguished from the dichotomous model?

Answer: Distinguishing the continuous model from the dichotomous model is one of these things that's going to be a long time in being developed. Finding a good test for that may never prove possible. I personally think that these two models are going to be accepted or rejected on the basis of their general fit to a wide set of data, rather than any one test. I'm not at all sure that the continuous model versus the dichotomous model can ever be distinguished by any simple experimental test. I think it's simply a matter of what you find more convenient, and which one works better over a wide set of experiments. A general test would be very hard.

If anybody has any idea on any firm tests to help distinguish those, please let me know. I rather doubt you can come up with one.

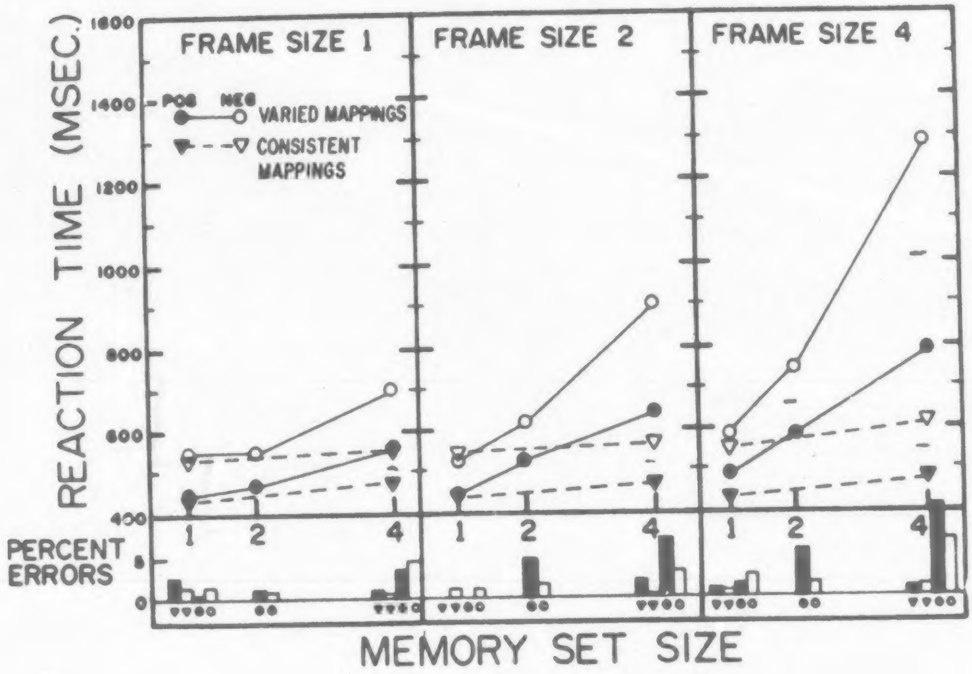
FRAME SIZE 1
FRAME SIZE 2
FRAME SIZE 4
POS
NEG
VARIED MAPPINGS
CONSISTENT MAPPINGS
REACTION TIME (MSEC.)
1600
1400
1200
1000
800
600
400
1
2
4
PERCENT ERRORS
5
0
MEMORY SET SIZE

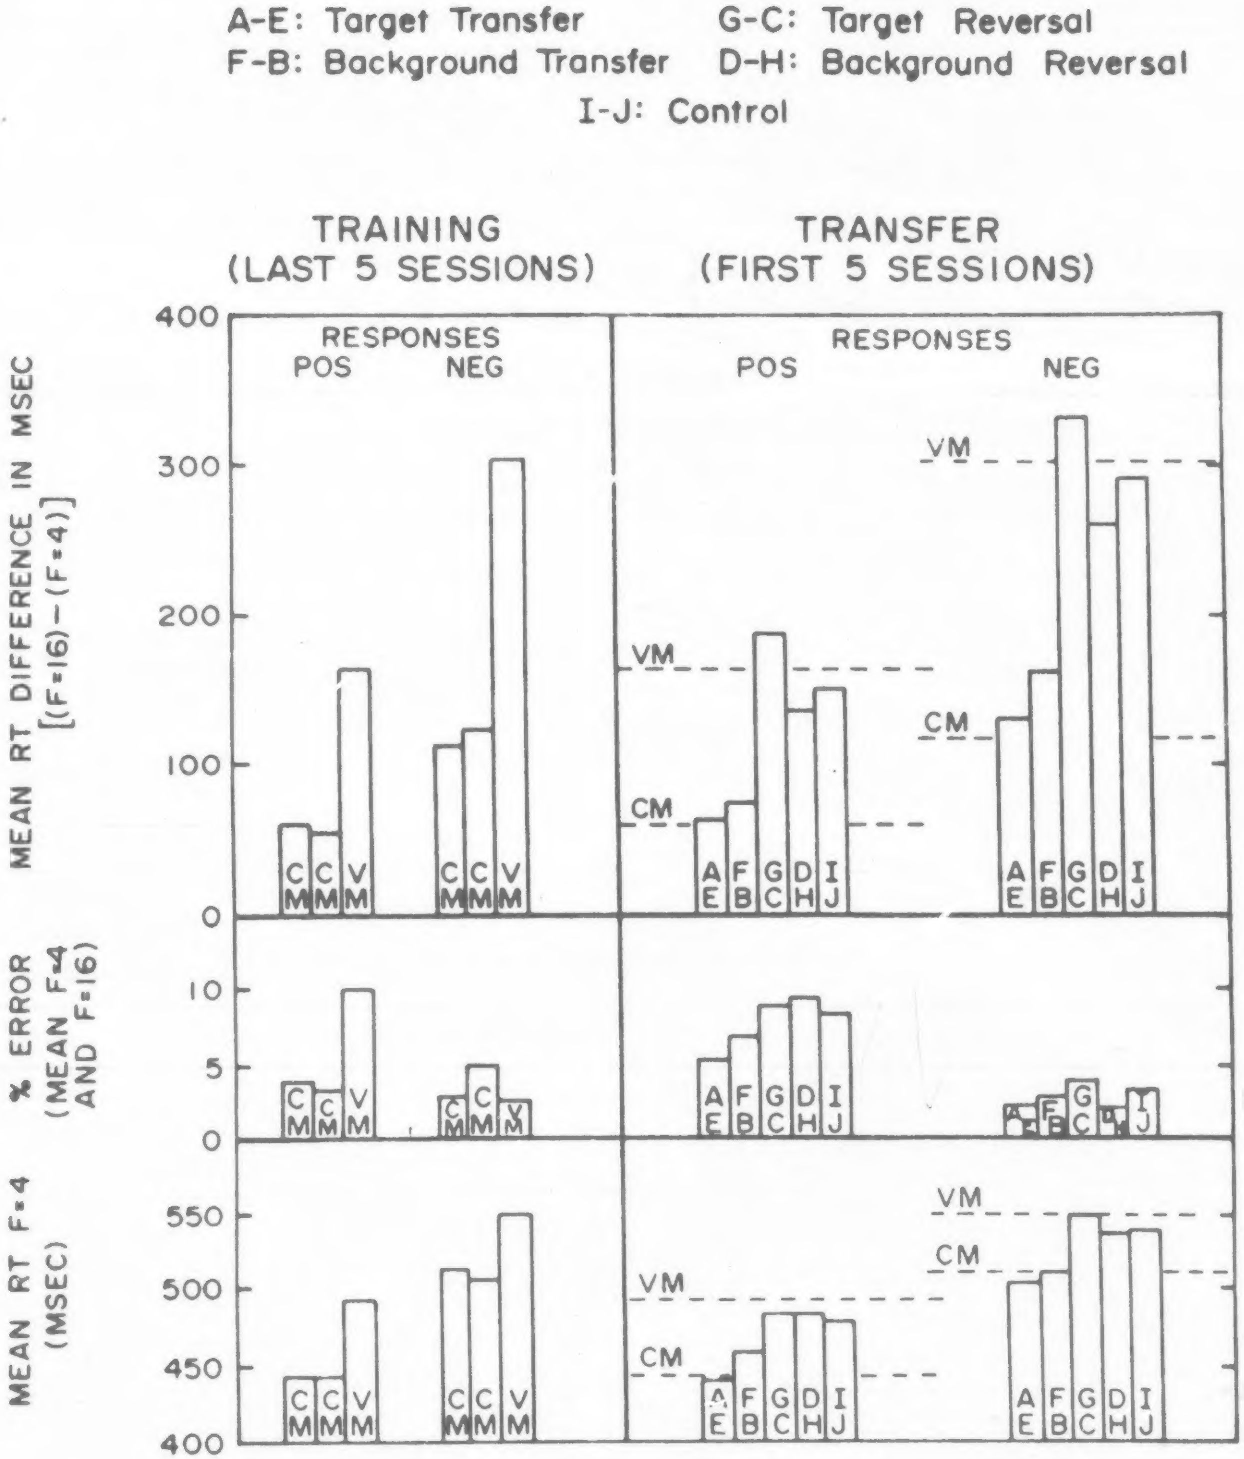
A-E: Target Transfer
F-B: Background Transfer
G-C: Target Reversal
D-H: Background Reversal
I-J: Control
TRAINING (LAST 5 SESSIONS)
TRANSFER (FIRST 5 SESSIONS)
MEAN RT DIFFERENCE IN MSEC [(F=16)-(F=4)]
% ERROR (MEAN F=4 AND F=16)
MEAN RT F=4 (MSEC)
RESPONSES
POS
NEG
VM
CM
CM
CM
VM
AE
FB
GC
DH
IJ
400
300
200
100
0
10
5
0
550
500
450
400

a

X X X 0 0 0
X X X 0 0 0
X X X 0 0 0
X X X 0 0 0
X X X 0 0 0

b

X 0 X 0 X 0
X 0 X 0 X 0
X 0 X 0 X 0
X 0 X 0 X 0
X 0 X 0 X 0

c

X X X X X
X X X X X
X X 0 X X
X X X X X
X X X X X